Reasons for Logic, Logic for Reasons

Pragmatics, Semantics, and Conceptual Roles

Ulf Hlobil and Robert Brandom

ROUTLEDGE

Reasons for Logic, Logic for Reasons

Reasons for Logic, Logic for Reasons presents a philosophical conception of logic—"logical expressivism"—according to which the role of logic is to make explicit reason relations, which are often neither monotonic nor transitive. This conception of logic reveals new and enlightening perspectives on inferential roles, sequent calculi, representation, truthmakers, and many extant logical theories.

The book shows how we can understand different metavocabularies as making explicit the same reason relations, namely normative-pragmatic, alethic-representational, logical, and "implication-space" metavocabularies. This includes a philosophical account of the pragmatic role of reason relations, treatments of nonmonotonic and nontransitive consequence relations in sequent calculi, a correspondence between these sequent calculi and variants of truthmaker theory, and the introduction of a novel kind of formal semantics that interprets sentences by assigning inferential roles to them. The book thus offers logical expressivists and semantic inferentialists new ways to understand logic, content, inferential roles, representation, and reason relations.

This book will appeal to researchers and graduate students who are interested in the philosophy of logic, in reasons and reasoning, in theories of meaning and content, or in nonmonotonic and nontransitive logics.

Ulf Hlobil is Associate Professor of Philosophy at Concordia University. His publications include "Limits of Abductivism About Logic" (2020) and *G.E.M. Anscombe: Aufsätze* (2014, with Katharina Nieswandt).

Robert B. Brandom is Distinguished Professor of Philosophy at the University of Pittsburgh. His books include *Making It Explicit* (1994), *Between Saying and Doing* (2008), and *A Spirit of Trust* (2019). He is a Fellow of the American Academy of Arts and Sciences and of the British Academy.

"Hlobil and Brandom join forces in this remarkable work of philosophy and logic, uncompromising in depth and in breadth. Their discursive rationalism provides the much-needed perspective of reasoners as they are engaged in critical reflection on their reasoning practices. We are invited to a bird's eye view of a rich network of logics that accounts for a variety of discursive practices and their relation to the world—and the sights are wonderful. *Reasons for Logic, Logic for Reasons* is a true masterpiece."

Gil Sagi, *University of Haifa, Israel*

Reasons for Logic, Logic for Reasons

Pragmatics, Semantics, and Conceptual Roles

Ulf Hlobil and Robert B. Brandom

Routledge
Taylor & Francis Group
NEW YORK AND LONDON

Cover image: © Sasa69M via Getty Images

First published 2025
by Routledge
605 Third Avenue, New York, NY 10158

and by Routledge
4 Park Square, Milton Park, Abingdon, Oxon, OX14 4RN

Routledge is an imprint of the Taylor & Francis Group, an informa business

© 2025 Ulf Hlobil and Robert B. Brandom

The right of Ulf Hlobil and Robert B. Brandom to be identified as authors of this work has been asserted in accordance with sections 77 and 78 of the Copyright, Designs and Patents Act 1988.

All rights reserved. No part of this book may be reprinted or reproduced or utilized in any form or by any electronic, mechanical, or other means, now known or hereafter invented, including photocopying and recording, or in any information storage or retrieval system, without permission in writing from the publishers.

Trademark notice: Product or corporate names may be trademarks or registered trademarks, and are used only for identification and explanation without intent to infringe.

ISBN: 978-1-032-36076-8 (hbk)
ISBN: 978-1-032-36077-5 (pbk)
ISBN: 978-1-003-33014-1 (ebk)

DOI: 10.4324/9781003330141

Typeset in Sabon
by Newgen Publishing UK

This work is dedicated to all the people who have thought so hard with us about these and related ideas over the years in the Research Group on Logical Expressivism (ROLE), especially Daniel S. Kaplan, Shuhei Shimamura, Ryan Simonelli, and Rea Golan—and to sets of pairs, structured by commutative monoids and a partition.

Contents

Preface *xi*

Introduction: Metavocabularies of Reason **1**
ROBERT B. BRANDOM

0.1 Two Traditions in the Philosophy of Language and Two Kinds of Metalanguage 1
0.2 Reason Relations, Vocabularies, and Metavocabularies 4
0.3 Intrinsic Metavocabularies for Reason Relations and Conceptual Contents: Logic 11
0.4 Implication-Space Intrinsic-Explicative Metavocabularies for Reason Relations and Conceptual Roles 17
0.5 Conclusion 25

1 A Pragmatic Metavocabulary **29**
ROBERT B. BRANDOM

1.1 Semantic and Pragmatic Construals of a Fundamental Bipolarity 29
1.2 Reasons and Reason Relations: Symmetries and Asymmetries 33
1.3 Bilateral Pragmatic Renderings of Reason Relations 42
1.4 Conclusion 52

2 Logical Expressivism and the Open Structure of Reasons **55**
ROBERT B. BRANDOM

2.1 A Pragmatic Definition of Implication and Incompatibility 55
2.2 Reasons and Logic: Logicism about Reasons 57
2.3 Reasons and Logic: Rational Expressivism about Logic 60
2.4 The Structure of Reason Relations 71
2.5 Even More Radically Open Structures of Reason Relations 81

2.6 Explicitation Closure and Rational Hysteresis 86
2.7 Conclusion: Expressivist Criteria of Adequacy for Logical Vocabulary 94

3 Introducing Logical Vocabulary: Making Reason Relations Logical 103
ULF HLOBIL

3.1 Sequent Calculi for Logical Expressivism 105
3.2 Making Reason Relations Explicit 121
3.3 Making Local Structure Explicit 131
3.4 Conclusion 142
3.5 Appendix 150

4 Truth-Taking and Truth-Making: How They Share a Rational Form 158
ULF HLOBIL

4.1 Rational Forms 159
4.2 Truth-Maker Theory 173
4.3 Articulating the Isomorphism 180
4.4 Accommodating Open Reason Relations 189
4.5 Conclusion 195
4.6 Appendix 200

5 Implication-Space Semantics: Content as Implicational Role 210
ULF HLOBIL

5.1 Formulating Implication-Space Semantics 211
5.2 Interpreting NMMS and Truth-Maker Theory 221
5.3 The Monoidal Structure of Reason Relations 228
5.4 Noncontractive Reason Relations 230
5.5 Implicational Role Inclusion 237
5.6 Conclusion 253
5.7 Appendix 256

6 Conclusion 281
ULF HLOBIL AND ROBERT B. BRANDOM

6.1 A Fourfold Perspective on Reason Relations 283
6.2 Rational Forms 290
6.3 Theorizing Open Reason Relations 293
6.4 The Logics of Open Reason Relations 297
6.5 Final Remarks 303

Epilogue: A Speculative Synthesis **308**
ULF HLOBIL

Rational Forms as Starting Points 309
Ideal Occurrences of Rational Forms: The World 314
The Pragmatic-Semantic Bifurcation Explained 319
Conclusion 323

Bibliography *327*
Index *332*

Preface

This book is the fruit of many years of collaboration, not merely between the two of us but rather a broader collaboration that included many colleagues, friends, and students. The Research Group on Logical Expressivism (ROLE) has met weekly for more than a decade, and many people have been members of the group during different time periods. The one whose contribution we most want to highlight and acknowledge specially in this preface is Daniel Kaplan, whose 2021 University of Pittsburgh Philosophy Ph.D. dissertation contains invaluable work and crucial results on the multiple conclusion sequent calculus and the implication-space semantics that we are presenting here.

Kaplan was the first in our group to realize how much things can be simplified and streamlined by moving from a single conclusion supra-intuitionistic sequent calculus to a multiple conclusion supra-classical sequent calculus. He realized that Ketonen's invertible connective definitions would make possible the kind of explicitation of a logic that in Chapter Three we call "explicitation by sequents" obeying "reflection," and the relations between the converse of this condition (which we call "representation" in Chapter Three) and the structural rule of Contraction. Perhaps Kaplan's most important contribution was the discovery of how Girard's phase-space semantics for linear logic can be adapted to our needs. Our version of the implication-space semantics that builds on this idea is presented in Chapter Five. Further, the notion of "implicational role inclusions" and its relations to strong Kleene, three-valued logics such as LP and K3 that are explored there is a development and integration into implication-space semantics of a formal apparatus he developed in our ROLE group for studying the relationships between the roles sentences have in consequence relations (known to us as his "right-arrow" machinery).

Our indebtedness to and reliance on Dan's earlier work on all these fronts is sufficiently substantial that our original plan was for him to serve as the principal author of the two chapters that build on constructions he

reports in his dissertation and in unpublished material circulated in the ROLE group. In the event, he realized that the press of difficult personal circumstances would make it impossible for him to do so to his standards and on a realistic time line, and gracefully withdrew from the project. As a result, though none of the particular sentences of the text presented in this work are specifically his, Dan's ideas are a pervasive influence on the reason relations that articulate the conceptual contents expressed by that vocabulary.

Other members of the ROLE group have also been important for the developments of the ideas that we are presenting here. Ryan Simonelli has done important work on how to think philosophically about the relation between logic and pragmatics, and he first saw some crucial connections between the ideas that we pursue in the ROLE group and certain issues regarding natural language conditionals, including connections between failures of monotonicity (MO) and failures of transitivity (Cut). Shuhei Shimamura has shown how quantifiers can be added to the logical systems that we present here, and he has done important work on invalidities and anti-sequents, as well as on various connections between our ideas and relevance logic and connexive logics. Rea Golan has also been instrumental in connecting the ideas that we pursued together with relevance logic, intuitionistic logic and other non-classical logics, as well as issues regarding meta-inferences.

In addition to the work by ROLE members and the weekly group discussions with them, we have been influenced by many friends and colleagues, including Jarsolav Peregrin, David Ripley, Viviane Fairbank, Lucas Rosenblatt, Katharina Nieswandt, Nabeel Hamid, Eduardo Barrio, Shawn Standefer, Federico Pailos, and Ladislav Koreň.

Work on the book was supported by the Czech EXPRO grant No. 20-05180X on inferentialism. Our work was also supported by a Canadian SSHRC Insight Grant. We are grateful for this support and for the help and comments we received from the members of the research teams funded by these grants.

Not only is this book deeply influenced and shaped by ideas of, and conversations with, colleagues and friends, but it is also a result of a collaboration between the two of us that started more than ten years ago. Nevertheless, we have decided to be explicit about who is the author of each chapter of this book, in the sense of having written the text of the chapters, although often on the basis of ideas that arose out of conversations and thinking about the issues together. Brandom is the chief author of the Introduction and Chapters One and Two. Chapters Three to Five as well as the Epilogue were written chiefly by Hlobil. We wrote Chapter Six together. Hence, we are chiefly responsible respectively for these parts of the text and any errors or otherwise problematic features that they might contain.

Introduction
Metavocabularies of Reason

Robert B. Brandom

0.1 Two Traditions in the Philosophy of Language and Two Kinds of Metalanguage

One of the principal distinguishing features of philosophical thought over the past one hundred and fifty years is its pursuit of the strategy of addressing philosophy's perennial concern with the nature of understanding or reason in general by investigating *language*, thought of as its essential, characteristic embodiment and expression. That guiding idea has been developed in two quite different directions, giving rise to two nearly independent and often antagonistic currents of thought. The first, dominant, better worked-out tradition focused on logic, and later, also formal semantics, as perspicuous mathematical metalanguages in which to articulate conceptual contents and their rational relations to one another. The other tradition focused rather on language as a kind of social practice, thought of as a prominent feature of the natural history of a certain kind of being: language as *parole* rather than *langue*, in de Saussure's terms.

The logico-semantic formalist tradition can be roughly delineated by tracing a central strand that leads from Gottlob Frege through Bertrand Russell and the Wittgenstein of the *Tractatus Logico-Philosophicus*, to Rudolf Carnap, Alfred Tarski, Saul Kripke, David Lewis, and Kit Fine. Its paradigm is the language of mathematics, from which it takes the model of justifying reasons as apodeictic proofs. The founding achievement of this tradition is Frege's invention of a logical calculus, itself a regimented mathematical vocabulary, adequate to express the reason relations and conceptual contents of arithmetic. That feat inspired the project of using the powerful expressive tools and explicative methods of the new logic to codify reasons and conceptual contents generally, with hitherto unapproachable precision and perspicuity. Beginning with Tarski's model theory (further developed, and packaged for and marketed to philosophers by Carnap and Quine), formal semantic metalanguages joined formal logical ones in the mathematical metaconceptual toolkit defining this

tradition. The expressive power of Tarskian model theory was substantially enhanced by its development in the form of possible worlds semantics, and then again in its further generalization into contemporary truth-maker semantics. In each case, the first application and proof of concept of the new semantic metalanguages was their demonstrable adequacy in providing a semantics for specifically *logical* vocabularies (in the case of possible worlds semantics, for modal logics). It was then shown how to extend those achievements to provide, first extensional (Tarski, Quine), then intensional (Lewis), and hyperintensional (Fine) semantics for *non*logical (and nonmathematical) expressions.

Avatars of the anthropological, natural-historical, broadly pragmatist tradition include the classical triumvirate of American pragmatists, Charles Peirce, William James, and John Dewey, and also the Heidegger of *Being and Time*, the Wittgenstein of the *Philosophical Investigations*, Richard Rorty, and such contemporary figures as Charles Travis and Huw Price. Where the formalist tradition is oriented by a conception of understanding and reasons as codified in artificial logical calculi and semantic metalanguages, the pragmatist tradition looks instead directly to natural languages, thought of as social practices and forms of life. In place of the exclusively monological character of reasoning as deriving, modeled on proof, characteristic of the other tradition, understanding shows up in this tradition as a social achievement, and reasoning as essentially dialogical: a matter of discursive practices of giving and asking for reasons, defending and challenging claims that amount to taking up positions in a contestable, public, normative space.

These two general philosophical approaches to language have been pursued to a remarkable extent independently of one another, rather than in conversation with each other. That fact both engenders and is supported by the impression that they are sufficiently antithetical that philosophers interested in language are obliged to choose between them, and to align themselves with one or the other camp. That unhappy impression is gainsaid by influential figures such as Davidson, Putnam, Sellars, Dummett, and Stalnaker, who, like Quine (and even Lewis), are distinguished by their labors in both traditions, and (unlike the time-slices of Wittgenstein) in developing and promoting their confluence. The two traditions ought by rights to be understood as focusing on different aspects of language: roughly, on the *meanings* of linguistic expressions, and on their *use*. In suitably broad senses, we might understand semantics as the study of meaning, and pragmatics as the study of use or discursive practices and abilities. So understood, semantics (even a semantics inspired by and paradigmatically applicable to logic) and pragmatics show up as complementary theoretical endeavors. The goal should be to synthesize semantic and pragmatic theories, so as better to understand the relations

between meaning and use, as two essential aspects of a single linguistic phenomenon.

That aspiration is compatible with acknowledging that there might still be good arguments for adopting a semantics-first or a pragmatics-first order of explanation. The thought that the norms that function as standards for assessments of the correct use of expressions must be sensitive to the meanings (for instance, the truth conditions) of the expressions used speaks for a semantics-first order of explanation. And the thought that all there is to confer meanings or contents on linguistic expressions is the (actual or proper) use of those expressions speaks for a pragmatics-first order of explanation. Perhaps the combination of those thoughts recommends rather a more balanced view that eschews claims of explanatory priority in favor of understanding each aspect as in principle intelligible only in terms of its relation to the other.

We pursue here a synthetic metalinguistic strategy. We think of semantics and pragmatics as theoretical undertakings pursued by talking about language in two different kinds of metalanguage: semantic metalanguages and pragmatic metalanguages (or, as we will say: metavocabularies). They make it possible to express, to make explicit, different aspects of a single phenomenon. Further, we think that powerful new insights into the nature of language can be provided by understanding that phenomenon as essentially, and not just accidentally, having aspects that can be made explicit by just the right semantic and pragmatic (and, indeed, logical) metavocabularies.

The lesson that emerges, we will argue, is a kind of discursive or linguistic *rationalism*. Language becomes visible as at base the medium of reasons, and reasoning as the beating heart of language. On the side of pragmatics, the fundamental speech act is that of making claims. The basic speech act of making claims, asserting, is to be understood in terms of practices of defending and challenging those claims, by making other claims that have the practical significance of giving reasons for and against them. Understanding claiming this way provides a path to understanding the claimable contents expressed by declarative sentences in terms of the role they play in relations of being a reason for or against—what we will call "reason relations." On the side of semantics, worldly represented states show up as what determine the reason relations of consequence and incompatibility that the sentences whose truth-makers and falsifiers they are stand in to one another: their roles in reason relations. By understanding the common topic that semantic and pragmatic metalanguages articulate aspects of, not just under the vague rubric of "language," but more specifically as the implicit reason relations that distinguish *discursive* practices as such, we can better understand not only the relations between the meaning and the use of linguistic expressions,

but also the relations between truth (the central concept of traditional semantics) and justification (the central concept of pragmatics, according to linguistic rationalism), in the form of practices of defending claims by giving reasons for them and challenging claims by giving reasons against them. One manifestation of that opposition between truth-based and justification-based approaches is the distinction between model-theoretic and proof-theoretic formal metavocabularies for codifying reason relations. We will have a lot to say about the relations between them in the rest of this book.

The multifaceted metalinguistic rationalism we present in the chapters that follow offers a new way of understanding why and how what is expressed by declarative sentences is both what can be semantically evaluated as true or false and what can be pragmatically accepted or rejected. It articulates a conception of sentential *conceptual content* in terms of the functional roles sentences play in virtue of standing to one another in reason relations of consequence and incompatibility, rational inclusion and exclusion, which show up in different, though intimately related guises, from the perspectives afforded not only by semantic and pragmatic metalanguages, but also by logical and metalogical ones. There is, of course, much more to meaning and to use than what is precipitated in conceptual content as we consider it. But seeing how the aspects of language that are addressed by logical, semantic, and pragmatic metalanguages can be understood as providing perspectives on the common topic of reason relations and roles with respect to them opens up avenues for reconciling and synthesizing the often apparently divergent things these vocabularies make it possible for us to say and think, about what we are saying and doing when we are talking and thinking.

0.2 Reason Relations, Vocabularies, and Metavocabularies

At the core of this book, then, is the rationalist explanatory strategy of understanding the nature of language in terms of what we call "reason relations." As addressed here, that is a genus with two principal species: implication and incompatibility. They correspond to being a reason for and being a reason against. One of our overarching aims is to offer a detailed, multifaceted account of reason relations as such. A subsidiary allied task is more precisely to delineate the object of study, by showing how reason relations appear from the quite different perspectives afforded by the vantage points of logic, semantics, pragmatics, and a new formal apparatus for articulating the conceptual roles declarative sentences play in virtue of standing to one another in relations of rational inclusion and exclusion. The common topic that shows up in various guises when we triangulate (really, quadrangulate) on it by considering what can be said about it in

these four disparate technical kinds of vocabulary is in many ways both familiar and recognizable, but also subtly yet importantly different from what is typically discussed under the heading of "reasons."

A closely related term of art is "vocabulary." We use it in a technical sense, to mean a lexicon or set of declarative sentences, together with an implication relation and an incompatibility relation defined on those sentences. To begin with, we can think of an implication relation as holding between a set of sentences that are its premises and a single sentence that is a conclusion that follows from, is a consequence of, or is implied by those premises. An incompatibility relation holds between a set of premises and a further sentence that those premises exclude, or rule out, or are incompatible with. The elements of the lexicon are repeatable sign designs of any sort. By calling them (declarative) "sentences" we just mean that they are what in the first instance stand to one another in reason relations of implication and incompatibility. (The reason relations among predicates—not considered in our treatment, which remains resolutely at the level of sentences—are derivative from these.) In virtue of standing to one another in reason relations of implication and incompatibility, what thereby count as declarative sentences express conceptual contents—specifically *propositional* contents. Those contents can be thought of as the functional roles the sentences play in constellations of implications and incompatibilities. Sentences express what can both serve as and stand in need of reasons (for and against, positive and negative).

This characterization is very abstract. Whatever plausibility there is to characterizing the elements of the lexicon of a vocabulary as "declarative sentences," and the roles they play as "conceptual or propositional contents" depends at least on the relations defining what we are calling a "vocabulary" really being reason relations of implication or consequence and incompatibility. (Even then, the connections being asserted between sentencehood and propositional contentfulness, on the one hand, and reason relations, on the other are substantive, potentially controversial, and so in need of specification, motivation, and justification. *Solvitur ambulando*.) According to this order of explanation, the key question is: what do we mean by talk of reason relations of implication and incompatibility? In virtue of what does something deserve to count as a consequence or incompatibility relation? On the approach we are pursuing, an answer to this question will ramify into answers to corresponding questions about sentences and conceptual content.

The idea is to identify reason relations in terms of the various vocabularies that can be used to specify them. Because these are vocabularies for talking about (the reason relations of) other vocabularies, they are *meta*vocabularies. Because it is in particular the reason relations of base vocabularies that they address, we can call them

rational metavocabularies—by contrast to metavocabularies that express other features of the lexicon: syntactic, phonological, or etymological metavocabularies, for instance. Two different strategies for explaining what reason relations are—corresponding to two different philosophical traditions of thought about language—are pursued in two different kinds of rational metavocabulary: semantic and pragmatic.

Semantic metavocabularies explain what is expressed by the reason relations of empirical base vocabularies by saying what features of the world those vocabularies capture in their consequence and incompatibility relations. The sentence "The coin is made of pure copper" implies the sentence "The coin would melt at 1085°C." because and in the sense that pure copper *necessarily* melts at 1085°C., and it is incompatible with the sentence "The coin is an electrical insulator," because it is *impossible* for pure copper not to conduct electricity. Semantic metavocabularies explain reason relations of implication and incompatibility by specifying what the sentences that stand in those relations *mean*, in the sense of how the world must be for what they say to be true. The sentences stand to one another in relations of implication and incompatibility because the objective states of affairs that are their semantic truth conditions stand to one another in modally robust relations of necessitation and noncompossibility. For sentences to be part of a vocabulary, and not just a lexicon, is for them to express meanings in the form of worldly truth conditions that are themselves related to one another by consequence and incompatibility relations that are specifiable in an alethic modal semantic metavocabulary.

Pragmatic metavocabularies explain what is expressed by reason relations of base vocabularies by saying what features of the discursive practices of *using* those sentences it is, in virtue of which practitioners count as practically taking or treating the sentences *as* standing to one another in relations of implication and incompatibility. Pragmatic metavocabularies make it possible to *say* what it is that language users *do* in virtue of which they are properly to be understood as practically taking or treating some sentences as implying others in the sense of taking assertion or acceptance of the premises as providing reasons for asserting or accepting the conclusions, and practically taking or treating some sentences as incompatible with others in the sense of taking assertion or acceptance of the premises as providing reasons against asserting or accepting the conclusions. Reason relations show up from the expressive perspective provided by pragmatic metavocabularies as normative standards for assessment of the correctness of rational defenses of and challenges to claims, made by offering other claims as reasons for or reasons against those claims. Using lexical items *that* way—as specified in the pragmatic metavocabulary—is using them *as* the sentences of a vocabulary.

These two kinds of metavocabulary are evidently different in many important ways. Semantic metavocabularies talk about the *meanings* expressed by the sentences of a base vocabulary, in the sense of how those sentences represent the objective world as being. Pragmatic metavocabularies talk about the *use* of sentences in virtue of which they express those meanings. As we will see later in much more detail, to do their job properly, semantic metavocabularies must use alethic modal vocabulary to make claims about what states and combinations of states of the world the base vocabulary talks about are and are not possible. To do their job properly, pragmatic metavocabularies must use deontic normative vocabulary to make claims about what acts, practical attitudes, and combinations of them are and are not appropriate, and what other acts and attitudes would and would not entitle an interlocutor to them. What can be said in alethic modal terms is substantially and importantly different from what can be said in deontic normative terms. The one concerns features of the objective world, the other features of the practice of discursive subjects. These are the two poles of the intentional nexus that links knowers and the known, minds and the world they understand and act in, representings and what is represented. We want to understand both kinds of thing, and the important relations between them. That is a philosophical challenge because of the substantial differences between them.

In claiming that both pragmatic and semantic metavocabularies can be regarded as *rational* metavocabularies, though, we are claiming that in spite of these differences of subject-matter, concern, and metaconceptual resources, they do have at least one common topic. Among the things they both discuss and seek to explain in their own terms are reason relations. This claim of an overlap in topics is a strong, and perhaps surprising one. Consider for the moment the example of incompatibility. The claim is that when the meaning-theorist understands the incompatibility of p and q as a matter of the fusion of any objective state that is a truth-maker of p with any truth-maker of q being an impossible state in an alethic modal sense, and the use-theorist understands the incompatibility of p and q as a matter of the position an interlocutor is in when accepting both p and q as being normatively out of bounds—one the interlocutor cannot be entitled to—while they are clearly using "incompatible" in different *senses*, there is a common *referent*, a "reason relation," that both are picking out and talking about, in their respective idioms. In the terms used above to introduce the idea of reason relations, we propose to understand the alethic modal semantic metavocabulary and the deontic normative pragmatic metavocabulary as offering different (meta)conceptual *perspectives* on a common object: the incompatibility of what is expressed by the declarative sentence p and what is expressed by

the declarative sentence *q*. Corresponding claims apply to reason relations of consequence or implication.

A sketch has been offered of how a kind of metalinguistic rationalism might serve to make mutually intelligible two traditions in the philosophy of language, to synthesize them by exhibiting them as articulating different dimensions of a unified phenomenon. The present point is that if the claim that it is possible to identify a rational structure common to what is expressed in pragmatic and semantic metavocabularies could be made out in detail, it would cast light on issues of much wider philosophical significance. For we can look at the relations between what is expressed in normative pragmatic and representational semantic metavocabularies in another way: as articulating the relations between the activities of talking and thinking, and what is being talked or thought *about*. This is the intentional nexus between subjects and objects, between mind and world, knowers and the known.

The American Pragmatists inherited from the German Idealists—who in turn inherited it from Romantic critics of the Enlightenment—the idea that the Cartesian tradition failed structurally, making itself a patsy for skepticism, by attempting to define subjects and objects independently of one another, and then later on facing the problem of how to bolt together things understood as having wholly disparate natures. (A distinction becomes a dualism when it is drawn in terms that make the relations between the distinguished items unintelligible.) The better strategy, they thought, was to start with a conception of intentionality as successful cognition (and action), and analyze the contribution of its subjective and objective aspects. One way to work out such a strategy begins with the thought that there is a kind of *structure* common to what normative pragmatic metavocabularies make it possible to say about the practices of discursive subjects using declarative sentences to manifest practical attitudes and undertake commitments, on the one hand, and what representational semantic metavocabularies make it possible to say about the modal relations among matter-of-factual states of the world those sentences come to represent by being so used, on the other.

The further thought that that common structure deserves to be understood as a specifically *conceptual* structure—in the sense of that term that indicates that the roles sentences that are used to make, defend, and challenge assertible and deniable claims, and what makes those claimables true or false, stand to one another in relations of consequence and incompatibility—begins to develop this strategy in the direction of a *conceptual realism*. It appeals to a sense of "conceptual" that does not restrict that term to the products of thought. For the worldly version of the relations that articulate the structure we are calling "conceptual" are relations of necessity and impossibility whose existence owes nothing to the

activities of discursive practitioners. They are objective relations, specified in the alethic modal vocabulary used to state laws of nature, and more generally, to specify subjunctively robust relations. The overall idea is that relations that *in some sense* correspond to those objective relations on the "subjective" side of the activities and practical attitudes of speakers and thinkers are reason relations of consequence and incompatibility specified in the deontic normative modal vocabulary of consequential commitments and preclusion of entitlements. Since the objective and subjective "aspects" of the claimed common conceptual structure are those specified in vocabularies using two different sorts of modality, alethic and deontic, respectively, this view can be denominated a specifically *bimodal* conceptual realism. In addition, the version of bimodal conceptual realism we will offer is *hylomorphic* in a recognizably Aristotelian sense. For it identifies a kind of *rational form* that is understood as common to thoughts and things, both to the assertibles/deniables commitment and entitlement to which is at issue in speech acts of claiming, and defending and challenging claims, and to the worldly states that make such claimables true or false. The common conceptual structure is a *rational* form because, as we will argue, the relations of consequence and incompatibility that show up in different guises in a whole constellation of intimately interrelated metavocabularies—not just pragmatic and semantic, but also logical and metalogical formal vocabularies dedicated to characterizing conceptual roles as such—are those that in the end underwrite practices of reason*ing*, by determining what is a reason for and against what.

One aim we are pursuing here, then, is to give a clear and definite sense to the term "conceptual realism," by means of a kind of metalinguistic functionalism. The genus of which the conceptual realism we explore is a species has been variously characterized by different philosophers. We take the view we develop to be a way of understanding what Frege means when he says "A fact is a thought that is true."[1] It is also one way of understanding the Tractarian claim that the world is the totality of facts, and that facts must be understood as what can be pictured in thought. So understood, it is a commitment Wittgenstein never relinquished, asserting in the *Investigations*: "When we say, and mean, that such-and-such is the case, we—and our meaning—do not stop anywhere short of the fact; but we mean: this—is—so" (Wittgenstein, *PI* §95). John McDowell (1996) explores the same sort of conceptual realist view in *Mind and World* under the slogan "The conceptual has no outer boundary."

These are deep waters. These pronouncements by great philosophers are mentioned to indicate that the stakes are high for the enterprise of explicating any form of conceptual realism. Here is a sketch of how we go about it. One of the key arguments we appeal to in filling in this neo-Aristotelian metalinguistic bimodal conceptual realism is a technical result

presented in Chapter Four. It articulates a clear and precise sense in which semantic and pragmatic metavocabularies of the kinds invoked here can be understood as offering different perspectives on a common topic. Greg Restall and David Ripley have worked out what they call a "bilateral" normative pragmatic understanding of the turnstile that marks implication relations in multisuccedent sequent calculi. They interpret the claim that a set of sentences Γ implies a set of sentences Δ—in the notation we will use here, written as the sequent $\Gamma \mathrel{|\!\sim} \Delta$—as meaning that the "position" one is in if one asserts or accepts all of Γ and denies or rejects all of Δ is normatively "out of bounds." (The account is called "bilateral" in virtue of treating the premise side of the turnstile differently from the conclusion side: taking the one to concern assertions and the other denials.) In the deontic normative vocabulary we will use to render this view in Chapter One, that constellation of *commitments* (to accept and reject) is one that one cannot be *entitled* to. The idea is that what discursive practitioners must *do* in order thereby practically to be taking or treating the sets of sentences Γ and Δ *as* standing to one another in the basic reason relation of implication is to treat the corresponding position as normatively inappropriate or out of bounds.

The Restall-Ripley bilateral normative pragmatic metavocabulary turns out to be related in surprising ways to what we take to be the most sophisticated contemporary heir of the representational metavocabularies of Tarskian model theory and later intensional semantics in terms of possible worlds (Lewis, out of Kripke, out of Carnap), namely Kit Fine's truth-maker semantic framework. It incorporates a flexible metaphysical vocabulary for specifying the objective universe that is available to be represented: it consists of "states," and all their mereological fusions. The states are divided into possible and impossible states. Possible worlds are rendered as mereologically maximal possible states. But there are also much more limited possible states, perhaps such as the frog being on the log or Pittsburgh being in Pennsylvania. In addition, there are many different impossible states, perhaps there being immortal mammals or faster-than-light particles. The representational content of declarative sentences is then understood in terms of assignments to them of sets of states as truth-makers and falsifiers. Global structural conditions on modally partitioned state spaces (for instance requiring that all the mereological parts of possibles states be possible) interact with conditions on assignments of truth-makers and falsifiers (for instance forbidding the truth-makers and falsifiers of logically atomic sentences to be overlapping sets).

We show below that if one defines semantic consequence in just the right way, a powerful, fruitful, and detailed isomorphism can be constructed relating truth-maker modal semantic metavocabularies and bilateral normative pragmatic metavocabularies.[2] The required definition

is that a set of sentences Γ semantically entails a set of sentences Δ just in case every fusion of truth-makers of all of Γ with falsifiers of all of Δ is an impossible state. (This is substantially different from any definition of consequence currently abroad in the truth-maker literature.) Even at this rough level of description, the analogy this establishes with the bilateral pragmatic definition of consequence is clear. For the Restall-Ripley normative pragmatics treats Δ as a consequence of Γ just in case the position one would be in by asserting all of Γ and denying all of Δ is out of bounds. Assertion and denial line up with truth and falsity, combinations of commitments (to accept and reject) in a position line up with fusion of truth-making and falsifying states, and normative out-of-boundness (preclusion of entitlement to the commitments incurred by those assertions and denials) of a compound practical position lines up with the modal impossibility of such a fusion state.

When Spinoza looked back on the relations between algebraic equations and geometric shapes on which Descartes modeled mind-world relations, he saw that the key feature distinguishing that new, more abstract notion of representation from earlier atomistic resemblance-based conceptions is the existence of a global isomorphism between the algebraic and geometrical vocabularies. Spinoza's slogan for the holistic insight that animated the representational revolution was "The order and connection of ideas is the same as the order and connection of things" (Spinoza, *Ethics* II, Prop. vii). The isomorphism between normative pragmatic and alethic representational semantic metavocabularies turns out to make possible in our setting a precise, tractable, and productive specification of that shared rational "order and connection." We think this is a good way to rationally reconstruct some central aspects of Aristotelian (and Scholastic) intelligible forms. This isomorphism is the core of our version of bimodal (deontic/alethic) metalinguistic conceptual realism.

0.3 Intrinsic Metavocabularies for Reason Relations and Conceptual Contents: Logic

The construction gestured at so far foreshadows an argument for understanding reason relations of consequence and incompatibility as constituting a structure common to representational meaning and to practical use, to truth-making and to justificatory practices, to the objective world talked about and to the activities of talking about it, to what is represented and to the representing of it. That these same reason relations show up from the two otherwise disparate perspectives afforded by (the right kind of) semantic and pragmatic metavocabularies offers some reason to think of those relations as central to language or discourse as such. It provides a metalinguistic functionalist path to what we have

called "linguistic rationalism." Such an approach is unusual, and so perhaps surprising in how it discerns rational forms amphibious between these different dimensions. For the sort of language philosophers in the analytic tradition are accustomed to think of as codifying "rational forms" is formal *logical* languages, rather than semantic or pragmatic metalanguages. We agree that logical vocabulary is indeed uniquely privileged among languages for expressing reasons. But it is different from the rational metavocabularies considered here in several important ways. Our conception of the expressive role characteristic of logical vocabulary is also different from received views. One can begin to get a sense of it by comparing and contrasting it with the metavocabularies already under discussion.

One important criterion of adequacy for both semantic and pragmatic metavocabularies as we understand them is that they offer expressive resources sufficient to provide *explanations* of the reason relations of arbitrary base vocabularies. They are able to say, each in its own distinctive idiom, both what it means for some sentences to stand to others in relations of implication or incompatibility and why they do: what further facts explain them standing in just those relations. Our preferred version of semantics offers, in effect, truth-makers for the claims that $\Gamma \# A$ (Γ is incompatible with A) and $\Gamma \mid\!\sim A$ (Γ implies A) in alethic modal terms of the impossibility of fusions of truth-making states of Γ with truth-making states of A, and truth-making states of Γ with falsifying states of A, respectively—that is, in terms of how the sentences of Γ and A represent the objective world to be. Our preferred version of pragmatics specifies how one must use sentences in order thereby to count as practically *taking* or *treating* them as standing in relations of implication or incompatibility. It does that in deontic normative terms of constellations of commitments to accept and reject the claimables they express being improper, inappropriate, or "out of bounds."

In order to make available substantive accounts of reason relations, these two very different kinds of rational metavocabularies bring to bear significant conceptual resources that are not already contained in the base vocabularies to which they apply: mereological and alethic modal metaphysical concepts and semantic interpretation functions from sentences to sets of truth-making and falsemaking states, and bilateral concepts of acceptance and rejection and the normative assessment of such practical attitudes, respectively. Because both of these kinds of metavocabulary appeal to conceptual resources beyond those intrinsic to the base vocabularies of which they are metavocabularies, and do so in service not just of characterizing the reason relations of those base vocabularies but of explaining them, the sorts of semantic and pragmatic

metavocabulary we consider can be denominated "extrinsic-explanatory" rational metavocabularies.

That is not the only kind of metavocabulary for specifying reason relations with which we are concerned. In addition to *extrinsic-explanatory* rational metavocabularies, there are also *intrinsic-explicative* ones. This latter kind of metavocabulary for reason relations restricts itself to the conceptual resources supplied by the base vocabularies whose reason relations it characterizes, and is used to make explicit those reason relations and the conceptual contents they articulate, rather than to explain why they are as they are, or what it is for them to be as they are. The principal phenomenon we initially seek to understand in these terms is *logic*. The first way logical vocabulary differs from the semantic and pragmatic metavocabularies considered so far is that it is an intrinsic, rather than an extrinsic metavocabulary for codifying reason relations. The rules by which logical vocabulary is introduced to extend any arbitrary nonlogical base vocabulary appeal to nothing more than the reason relations sentences of the base vocabulary stand in to one another.

To appreciate the second basic feature distinguishing the specifically logical codification of reason relations, one must think about the metavocabulary by which logical vocabulary is introduced. There are many vocabularies that can do this, including Hilbert-style axiomatics and Tarskian model theory. We use a version of Gentzen's sequent-calculus approach. Gentzen's basic innovation was to treat reason relations, paradigmatically implications, as objects, called "sequents," that can be referred to and manipulated, and their metainferential relations made explicit in a mathematical metavocabulary. The sequent-calculus metavocabulary can be thought of as applying to an arbitrary nonlogical base vocabulary (that is, a vocabulary whose lexicon consists of logically atomic sentences). The sequent "$\Gamma \mid\!\sim \Delta$" says that in the base vocabulary, the set of sentences Γ implies the set of sentences Δ. (The sequent "$\Gamma, A \mid\!\sim$ " says that the sentence A of the base vocabulary is incompatible with the set of sentences Γ. The empty right-hand side of the turnstile says that the set of sentences on the left-hand side is materially incoherent.) Metainferential relations among sets of implications/incompatibilities of the base vocabulary, in the form of sequents, can then be expressed with horizontal lines, which are understood as saying that whenever all the sequents above the line are good (the reason relations they codify hold in the base vocabulary), then so are all the sequents below the line. In this way, structural metainferential relations among sequents can be expressed. For instance,

$$\frac{\Gamma \mid\!\sim \Delta}{\Gamma, A \mid\!\sim \Delta}$$

says that the top sequent can be weakened on the left by the sentence A, in the sense that adding A as a further premise to $\Gamma \mid\!\sim \Delta$ will not turn that good implication into a bad one. If "Γ", "Δ", and "A" are read as schematic letters standing for arbitrary sets and sentences, rather than as constants standing for particular ones, this formula says that the reason relations of the base vocabulary is structurally monotonic, in the sense that adding premises never infirms implications. This sequent-calculus metavocabulary allows for efficient expression of the reason relations that hold in any arbitrary base vocabulary, including metainferential relations. But it is essentially just a notation, requiring no substantial additional conceptual resources beyond what is provided by the base vocabulary whose nonlogical implications and incompatibilities it specifies explicitly.

Perhaps surprisingly, the spare sequent-calculus notation (a version of which we have already been using informally in this *Introduction*) turns out to be sufficient to formulate rules for adding logical vocabulary to any arbitrary base vocabulary, and (most importantly) computing the reason relations of the extended vocabulary from those of the base. The idea is first to extend the lexicon of the base vocabulary, by syntactic rules that specify that the base lexicon is included in the logically extended lexicon, and that if A and B are sentences of the extended lexicon, then so are A, $A \to B$, $A \wedge B$, and $A \vee B$. The metainferential connective rules then specify the reason relations that (sets of) sentences of the logically extended lexicon stand in to one another, entirely in terms of sequents codifying the reason relations that sentences of the base lexicon stand in. These include rules such as:

$$\frac{\Gamma, A \mid\!\sim}{\Gamma \mid\!\sim \neg A}$$

and

$$\frac{\Gamma, A \mid\!\sim B}{\Gamma \mid\!\sim A \to B}$$

The complete logically extended vocabulary (lexicon of logically complex sentences plus reason relations among them) can then be computed from the base vocabulary (lexicon of logically atomic sentences plus reason relations among them). We say that a corresponding logically extended vocabulary can be *elaborated from* any arbitrary base vocabulary. Implications and incompatibilities (and metainferences involving them) that hold in every logical extension of a base vocabulary, no matter what base vocabulary it is elaborated from, can then be said to hold in virtue of logic alone. These are sequents such as $\Gamma, A \to B, A \mid\!\sim B$ that can be shown to hold no matter what implications hold in the base vocabulary, simply

in virtue of the rules for determining the reason relations that logically complex sentences stand in (in the context of the basic structural principle of Containment—the principle that among the conclusions implied by any premise-set are the members of that premise-set).

The sequent-calculus vocabulary is accordingly a rational meta-vocabulary—a vocabulary for specifying the reason relations of some other vocabulary—that has the special feature that it permits the elaboration of arbitrary base vocabularies into vocabularies over lexicons that extend the lexicons of the base vocabularies by adding logically complex sentences formed by combining the sentences of the base vocabulary with logical operators. Rules for those operators formulated in the sequent-calculus vocabulary conservatively extend the reason relations of the base vocabulary, in the sense that the implications and incompatibilities that hold among logically atomic sentences in the logically extended vocabulary are just those that already held among them in the base vocabulary. And the connective rules formulated in the sequent-calculus vocabulary do this while appealing to no resources outside of those provided already by the reason relations of the base vocabularies to which they apply and which they conservatively extend.

That is, sequent-calculus metavocabularies are *intrinsic* rational metavocabularies of a distinct kind. What distinguishes them from other intrinsic rational metavocabularies (examples of which will be introduced just ahead) is their serving to extend arbitrary base vocabularies by adding to them a further sort of vocabulary: logical vocabulary. And they elaborate all the reason relations of the extended vocabulary solely from the reason relations of the base vocabulary to which they are applied. This is no doubt a remarkable achievement. But one might nevertheless want to know: what is it good for? Why do this? The answer to this question provides the third, most specific characterization of what distinguishes *logical* vocabularies from all other kinds.

The idea we pursue is that logical vocabulary is distinguished above all by the characteristic *expressive* role that it plays with respect to the reason relations of the base vocabulary to which it is added. When the reason relations of the logically extended vocabulary are suitably elaborated from those of a base vocabulary, it becomes possible for the first time to *say explicitly*, in the extended vocabulary, what implications and incompatibilities hold in that base, and also in its logical extension. Logical vocabulary makes explicit the reason relations of *any* nonlogical vocabulary to which it is added and from which its own reason relations are elaborated, paradigmatically by rules expressed in a sequent-calculus metavocabulary. When introduced by suitable sequent rules, a set of sentences Γ implies the conditional $A \rightarrow B$ just in case Γ together with A implies B, and Γ implies the negation $\neg A$ just in case A is incompatible

with Γ. (This is what the sequent rules for the conditional and negation displayed above stipulate.) Logically extending a base vocabulary makes it possible to *say*, in the resulting vocabulary, what the reason relations of that base vocabulary are (and as a bonus, also those of the logically extended vocabulary elaborated from it). In addition to these proper, first class, *expressive* sentential connectives, there are also auxiliary, second class, merely *aggregative* connectives of conjunction and disjunction.

Logical vocabulary, we say, is *el*aborated from and *ex*plicative of (LX for) the reason relations of arbitrary base vocabularies. Because it can extend any and every base vocabulary (that meets the most minimal conditions), the ideal logic is *universally* LX. Logical vocabulary can be demarcated by this distinctive *explicative* role. The constellation of the sequent calculus metavocabulary and the logical vocabulary it introduces stands in an intrinsic-explicative metalinguistic relation to the reason relations of any base vocabulary whatsoever. In general, it is not hard to arrange that the reason relations of some extended language can be computed entirely from the reason relations of any arbitrary base vocabulary: that the extended reason relations can be *elaborated from* those of the base. Care is needed in formulating the rules that introduce the logical connectives, however, in order to see to it that they can perform their full *explicative* function. The rules of the logics we propose can be shown to be *expressively complete* in a strong sense. For any arbitrary set of base implications and incompatibilities, it is possible to compute an implication (or a set of implications for some of the systems below) in the logically extended vocabulary that is derivable just in case exactly those reason relations in the base vocabulary hold. And for any set of implications-and-incompatibilities in the logically extended vocabulary, there is a set of reason relations in the base vocabulary that holds just in case those logically complex reason relations are derivable.

It is even harder to formulate a set of definition rules for sentential logical connectives that is *universally* LX, in the sense that a logically extended vocabulary can be elaborated from and will then be explicative of *any and every* base vocabulary. In particular, almost all extant logics either presuppose that the base vocabularies they extend satisfy strong global structural constraints—paradigmatically the monotonicity and transitivity at the core of traditional understandings of specifically logical consequence as a kind of closure operator—or retroactively impose some such global structure, thereby failing to be conservative over some substructural base vocabularies. While we believe that purely *logical* consequence does have a global closure structure (and that logical inconsistency is monotonic), we argue that this is not in general true of *non*logical reason relations, and so is not and cannot be true of the full consequence (and incompatibility) relations of logical extensions of arbitrary base vocabularies.

The philosophy of logic that takes the ideal logical vocabulary to be one that is universally LX—that the reason relations of the logically extended vocabulary can be elaborated from and are explicative of any arbitrary base vocabulary to which the logical connective rules are applied—we call "logical expressivism." We show how to define logics that fully satisfy this ideal, and then how their capacity to express reason relations can be further augmented to include local metainferential structural properties of arbitrary base vocabularies. Our favorite (due to Kaplan (2018) and Hlobil (2018)), which we call "NMMS," for "NonMonotonic, MultiSuccedent" logic, is expressively complete for arbitrary substructural material base vocabularies, including nonmonotonic and nontransitive ones. It is supraclassical when applied to base vocabularies meeting the most minimal structural condition (Containment, which requires that every premise-set implies all the premises it explicitly contains), and it determines exactly the consequence relation of classical logic when applied to base vocabularies that consist entirely of instances of Containment. (There is also a supraintuitionist version that converges on the consequence relation of intuitionist logic under the same circumstances (Hlobil, 2016).) We call these and the various extensions of them that we define "expressivist logics."

0.4 Implication-Space Intrinsic-Explicative Metavocabularies for Reason Relations and Conceptual Roles

Intrinsic-explicative rational metavocabularies, of which logical vocabularies are a paradigm, are philosophically important because they specify the "rational forms" of the reason relations that the above-mentioned isomorphism shows are common to the richer, more committive specifications of those reason relations provided by semantic and pragmatic metavocabularies. Characterizing those common rational forms requires abstracting away from the substantive concepts *extrinsic* to the base vocabularies for which they are metavocabularies that permit semantic and pragmatic *explanations* of reason relations. If the devil is in the (extrinsic, semantic and pragmatic) details, then the angels will be found working with rational forms at the more abstract (intrinsic) level of reason relations alone. Logical vocabularies make reason relations explicit in terms that appeal only to the conceptual resources supplied by the base vocabularies from which they are conservatively elaborated. They are in that sense *intrinsic* vocabularies for specifying reason relations. Logical vocabularies, however, are not purely *meta*vocabularies, in the sense in which semantic and pragmatic rational metavocabularies are. The sequent-calculus vocabularies in which we say how to elaborate arbitrary base vocabularies into logically extended vocabularies with the capacity to codify reason relations

are genuine metavocabularies in that sense. Like semantic and pragmatic metavocabularies, they do not extend the base vocabularies for which they are metavocabularies. They are *purely* metalinguistic, talking *about* expressions in the base vocabulary, rather than *using* them. But the new, logically complex sentences introduced by adding logical vocabulary to base vocabularies by rules formulated in the sequent-calculus metavocabulary do use the expressions of the base vocabulary, rather than merely mentioning them. Logical vocabulary is a hybrid or mongrel kind of metavocabulary. It plays the expressive role of explicitating reason relations: making them explicit, constructing sentences intelligible as *saying that* relations of implication and incompatibility hold. That is a broadly metalinguistic expressive function. But logical vocabulary performs that explicative expressive function by using the sentences whose reason relations it articulates, rather than by talking *about* them (mentioning them).

These observations raise the question of whether there is a purely intrinsic-explicative vocabulary for specifying reason relations that is a rational *meta*vocabulary in the sense of being genuinely and wholly metalinguistic. The answer is "yes," and seeing how such a metavocabulary works and what it can do illuminates both the genus of intrinsic-explicative vocabularies for reason relations, which it shares with the more familiar logical species, and the concept of reason relations (and so vocabulary) as such. Our candidate, informed by work due to Daniel Kaplan (2022), is an implication-space metavocabulary for specifying both reason relations and the conceptual roles sentences play in virtue of standing to one another in such reason relations. Very roughly, where Gentzen's sequent-calculus metavocabulary treats implications as basic objects in a proof-theoretic formalism, Kaplan's implication-space metavocabulary treats them as basic objects in a model-theoretic formalism. It represents the current state of the art in inferentialist semantics.

Inferentialists have long thought that the universe from which semantic interpret*ants* are drawn or from which those interpretants are built—the analogue of the universe of mereologically structured worldly states out of which semantic interpretants (propositions) as pairs of sets of truth-making states and falsifying states are built—should consist of implications (including incompatibilities coded as implications) and sets of them. Until now, however, inferentialists had assumed that the items that are in the first instance semantically interpret*ed* should be the same declarative sentences that stand to one another in reason relations of implication and incompatibility. This methodological commitment was thought to be controversial only from below, inasmuch as some representationalists in the model-theoretic tradition begin instead with singular terms and predicates, building sentences out of subsentential parts, thought of as more basic. By contrast, inferentialists understand such components as emerging only as

the result of analysis of sentences by assimilating sentences into equivalence classes according to which lexical-syntactic substitution relations among them preserve the goodness of implications among them (see Brandom, 1994, Chapter Six).

Kaplan's (2022) first conceptual innovation was the idea that thoroughgoing inferentialists ought to treat the most basic units being interpreted, no less than the semantic intepretants assigned to them, as being implications, rather than the sentences that make up their premises and conclusions. Only at a second, subsequent stage, would semantic interpretation be extended from implications to the sentences they contain. He accordingly begins with a universe of candidate implications, together with a partition of that universe into a distinguished set of *good* implications—ones whose conclusions really follow from their premises—and the rest. This universe of candidate implications with a distinguished subset is an *implication space*. Given a lexicon of sentences, it is just the space consisting of all the ordered pairs of sets of sentences of that lexicon. Thought of as the space from which semantic interpret*ants* are drawn, such an implication space is analogous to Fine's space of states, partitioned into a distinguished set of possible states, and a residue of impossible states. Any base vocabulary determines such an implication space, since the lexicon of the vocabulary suffices to define the points (candidate implications as ordered pairs of sets of sentences of the lexicon), and the reason relations of the vocabulary suffice to determine the distinguished set of good implications. Since that is all that implication-space semantics requires to define (at the second stage of semantic interpretation) the implicational roles played by sentences, it will be an intrinsic semantics, appealing only to the conceptual resources afforded by the base vocabularies for which it offers a semantic interpretation.

The new idea is to begin the process that will end in semantic interpretation of sentences by using the points of the implication space generated by any arbitrary base vocabulary—the candidate implications, both good and bad—also as the source of the items to be semantically interpreted. What could count as a semantic interpretation of an implication? What would make one think the idea even makes sense? Here is a rationale. We are exploring the idea of understanding meaning to begin with in terms of *reasons* instead of understanding it in terms of *truth*. That is to understand meaning in terms of a dyadic *relation* (between sets of sentences) instead of in terms of a monadic *property* (of sentences). On the approach that takes truth as basic, one starts with assignments to sentences of a truth *value*: as true or false, correct or incorrect, good or bad (*as* a representation). However, although assignments of truth values are the beginning of semantic interpretation on this approach, they are not the end. To get a notion of meaning that corresponds to what one grasps (however

imperfectly) when one understands a sentence, one must advance from consideration of truth *values* to consideration of truth *conditions*. (One must add to a semantic conception of the Fregean *Bedeutung* of a sentence a semantic conception of its Fregean *Sinn*.)

Thought of as a universe of potentially semantically evaluable items, the implication space determined by any given base vocabulary has a natural analogue to truth values, but on the side of reasons. For reason relations of implication and incompatibility (represented as pairs of sets of sentences of the lexicon) are divided into good ones and bad ones, ones that hold and ones that do not, correct implications (codifying real reasons) and incorrect ones (codifying possible but not actual reasons). At the extensional semantic ground level, one can say *that* a sentence is true, and in the reason-based setting one correspondingly can say at the extensional semantic ground level *that* an implication is good or an incompatibility holds. Given that analogy, the question becomes: what stands to implication (reason relation) *values* (good/not-good) as truth *conditions* stand to truth *values*?

The idea behind truth conditions (and Fine's generalization to truth-makers and falsifiers) is that apart from the question of whether a truth-candidate actually *is* true or false, there is the question of what it would take to *make* it true—what things would have to be like for it to count as correct in this distinctive semantic sense. The idea behind the first stage of implication-space semantics is that apart from the question of whether a candidate implication actually *is* good (according to the partition of the space of candidate implications into good and bad determined by the underlying base vocabulary), there is the question of what it would take to *make* it good. In the special case of reason relations that already do hold, candidate implications that are good, this takes the form of asking about the circumstances under which it would *remain* good. That is the *range of subjunctive robustness* of the implication.

To make this thought slightly more definite, we can consider a single-succedent sequent $\Gamma \mathrel{|\!\sim} A$. Suppose it is a good implication—that A actually follows from Γ according to the reason relations of the base vocabulary. We can think of the range of subjunctive robustness of this implication as the set of additional premises that could be added to Γ without turning it from a good to a bad implication. The range of subjunctive robustness of $\Gamma \mathrel{|\!\sim} A$ would then consist of all the good implications of the form $\Gamma, \Delta \mathrel{|\!\sim} A$. If $\Theta \mathrel{|\!\sim} A$ does *not* hold (is not in the distinguished set of good candidate implications of that particular implication space) then its range of subjunctive robustness is the set of all good implications that result from adding further premises that would make it good: all the $\Theta, \Delta \mathrel{|\!\sim} A$ that *do* hold. Those sets represent what additional premises are required to make the original candidate implication into one that is good. The range of

subjunctive robustness of a candidate implication is its semantic intepretant in the form of its *good-makers*, as in Fine's truth-based semantic setting the semantic interpretants of sentences are their *truth-makers* (and falsifiers).

Grasping ranges of subjunctive robustness in this sense is an essential part of understanding reason relations in ordinary vocabularies. Working with ordinary, defeasible nonlogical implications, one might think that the conclusion:

C) The lioness will attack the gazelle.

does not follow from

P1) The lioness can see the nearby wounded gazelle.

but that it does follow from

P2) The lioness is healthy.
The lioness is hungry.
The lioness can see the nearby wounded gazelle.

That is understanding something about what it would take to make $P1 \mid\!\sim C$ a good implication, even though as it stands it is not. $P2$ is adding premises that yield a good implication. And if one thinks further that C would follow from

P3) The lioness is healthy.
The lioness is hungry.
The lioness can see the nearby wounded gazelle.
Today is Tuesday.

but would not follow from

P4) The lioness is healthy.
The lioness is hungry.
The lioness can see the nearby wounded gazelle.
The lioness will immediately be struck by lightning.

then one appreciates something important about the range of subjunctive robustness of the implication $P2 \mid\!\sim C$.

The ranges of subjunctive robustness of candidate implications are their "goodness" conditions, as truth conditions are the "goodness" conditions of sentences. For an implication to be good in the reasons-first semantic setting is for its premises to provide reasons for its conclusion

(or reasons against, in the case of incompatibilities), while for a sentence to be good in the truth-first semantic setting is for it to be true. The advance from a conception of semantic goodness to a conception of meaning is the advance to consideration of circumstances under which a reason relation or sentence *would* be good. In the truth-maker/falsifier semantic setting, the circumstances are sets of worldly states, specified in some vocabulary that is not in general part of the base vocabulary whose sentences are being semantically interpreted. In the implication-space semantic setting, the circumstances are additional premises (and, in the fully general multisuccedent case also additional conclusions) that would make or keep the reason relation good. By contrast to the truth-maker setting, in the implication-space setting, those further premises and conclusions are just more sentences of the lexicon of the base vocabulary. That is why implication-space semantics counts as *intrinsic*.

Nothing has yet been said about why or how implication-space constructions can be *explicative* of the conceptual contents conferred on sentences in the lexicon by the reason relations of the base vocabulary, from which the implication space is derived. That expressive achievement is the task of the second stage of semantic interpretation in the implication-space setting. During that second stage, semantic interpretants for sentences (in the form of conceptual roles) are computed from the semantic interpretants of implications (their ranges of subjunctive robustness) in which those sentences appear as premises or conclusions. Essentially, the conceptual role of a sentence is represented as the ordered pair of the set of items with the same range of subjunctive robustness as the bare occurrence of the sentence as a premise and the set of items with the same range of subjunctive robustness as the bare occurrence of the sentence as a conclusion. It turns out that conceptual roles of this form perspicuously and tractably codify the contributions the occurrence of a sentence makes to the goodness of implications and incompatibilities in which it occurs as premise or conclusion.

One measure of the extent to which the implication-space specification of the conceptual roles sentences play in virtue of standing to other sentences in the reason relations they do is an expressively powerful one is that various semantically significant relations among sentences correspond to natural operations on conceptual roles specified in these terms. To begin with, the sequent-calculus rules that extend arbitrary base vocabularies by adding sentential logical connectives all correspond to simple ways of combining implication-space conceptual roles. In this way, a model-theoretic inferentialist semantics becomes available that is sound and complete for the aforementioned expressive logic NMMS. The implication-space semantics shows how to compute the conceptual roles of arbitrary logically complex sentences from the conceptual roles

of logically atomic sentences of any base vocabulary—even when the base vocabulary, and so its (conservative) logical extension, are radically substructural, including those that do not satisfy the metainferential structural closure conditions of monotonicity and transitivity. To do this, the implication-space rational metavocabulary must make explicit the conceptual roles played by sentences of all those base vocabularies, as well as their logical extensions. It is universally explicative of sentential conceptual roles. And since implication spaces can be constructed using no resources other than those supplied by the spare specifications of arbitrary, even substructural base vocabularies—just sets of sentences and set-theoretic constructions from them representing their reason relations—the implication-space model-theoretic semantics qualifies as a *universal* intrinsic-explicative rational metavocabulary.

One reason to think that the implication-space metavocabulary for explicating the conceptual roles sentences play in virtue of standing to one another in the reason relations that define any base vocabulary cuts at important joints is that we can erect on top of it a calculus codifying metainferential relations among those conceptual roles (see also Kaplan, 2022). Metainferences of various kinds can be defined precisely, systematic combinations of them recursively constructed, and the effects of those combinations computed. The result is a principled botanization of constellations of metainferences that offers revealing characterizations of a number of logics that have been the subject of intense interest among logicians and philosophers of logic over the past few decades. The trilogics, paracomplete K3 (Strong Kleene) and its dual, paraconsistent LP (Logic of Paradox), for instance, show up as the logics of premissory and conclusory metainferences, respectively—in virtue of the same features that make them show up as logics of truth-value gaps and gluts, respectively, in the truth-based semantic setting. The nontransitive classical logic ST (Strict-Tolerant) and the closely related logic TS (Tolerant-Strict) fall into place neatly in relation to them. In treating metainferential relations among conceptual roles as objects that can be combined and manipulated, this calculus stands to conceptual roles as the sequent calculus stands to the sentences that are the relata of the implication relations it codifies as sequents. This intrinsic rational metametavocabulary, built on top of the implication-space inferentialist model-theoretic semantics for conceptual roles, provides the expressive power to make explicit a hitherto unexplored level of metainferential reason relations among those roles, and thereby offers an illuminating new semantic perspective on the relations among a variety of well-studied logics.

The implication-space metavocabulary provides a model-theoretic semantics for the conceptual roles sentences play in virtue of standing to one another in reason relations of implication and incompatibility.

It is a reason-based inferentialist semantics, rather than a truth-based representational semantics like truth-maker semantics. By contrast to the proof-theoretic treatment of reason relations by the sequent calculus, the implication-space metavocabulary assigns sets of implications as the semantic interpretants of implications, and set-theoretic constructions out of those sets as the semantic interpretants of sentences, and then operates on and manipulates those semantic interpretants to codify reason relations and conceptual roles. In fact, it does so in a way that can be shown to be isomorphic to the truth-maker model-theoretic semantics. Each constructs bilateral semantic interpretants for sentences (truth-makers and falsifiers in the one case and premissory and conclusory roles in the other) from a universe consisting of a set of points (states and candidate implications, respectively) with a distinguished subset (possible states and good implications). In both cases, the universe is taken to be structured by a commutative monoid (fusion of states and a corresponding operation combining candidate implications according to their ranges of subjunctive robustness). Nonetheless, the implication-space metavocabulary provides an intrinsic semantics, since it appeals to nothing that is not made available by the base vocabulary to which it is applied: sets of sentences and their reason relations. Implication-space semantics is something like the *intrinsification* of truth-maker semantics—in a way formally analogous to, but more expressively powerful than, Fine's use of intrinsic "canonical models."

When this structural isomorphism of implication-space and truth-maker semantics—which holds between the universes from which semantic interpretants are drawn, the interpretants themselves, and the way reason relations of consequence and incompatibility are determined for sentences in terms of their semantic interpretants—is appreciated in detail, and considered in context with the orthogonal isomorphism at the level of reason relations between the truth-maker alethic modal semantic metavocabulary and the deontic normative bilateral pragmatic metavocabulary, it becomes clear that the implication-space semantics makes explicit the abstract rational forms common to those two extrinsic-explanatory metavocabularies of meaning and use. Those rational forms are just the conceptual roles the implication-space semantics characterizes. As an intrinsic-explicative rational metavocabulary, the implication-space explication of those conceptual roles abstracts away from the extrinsic appeals to the possibility or impossibility of fusions of worldly states and the propriety or impropriety of constellations of practical commitments to accept and reject that the truth-maker semantic and bilateral pragmatic metavocabularies call on to explain implication and incompatibility relations. As the pure (intrinsic) metavocabulary of conceptual roles, inferentialist implication-space semantics is our candidate for the language of the angels, who think by manipulating the abstract rational forms

shared by representations of worldly states on the objective side of the intentional nexus and manifestations of normative practical attitudes on the subjective side. We are pleased and grateful to report that the intrinsic-explicative metavocabularies of implication-space semantics and logic can also be used by embodied fallen creatures to achieve that angelic rational self-consciousness.

0.5 Conclusion

The topic of this book is metavocabularies of reason, or rational metavocabularies. "Vocabulary" is a term of art. We mean by it a set of items called "sentences," (a "lexicon") and a set of relations on (sets of) those sentences that we call "reason relations." The justification for calling the elements of the lexicon "sentences" is that they stand to one another in reason relations. So the question becomes: what makes a set of relations deserve to count as "reason" relations in a sense that in turn justifies calling what stands in such relations "sentences"?

As we use the term, there are two principal kinds of reason relation: "implications" and "incompatibilities." Those indeed sound like varieties of rational relation—but clearly the issue then becomes what qualifies relations as those of implication or consequence (a kind of rational inclusion) and incompatibility or inconsistency (a kind of rational exclusion). For the last century or so, logicians have offered responses to this question that look to the *structure* of the relations in question. Inspired by structural features of the best-behaved logics of their day, Tarski and Gentzen (founders of the model-theoretic and proof-theoretic metavocabularies for logic) agreed that consequence relations are reflexive, monotonic, and transitive, incompatibility relations are monotonic, and that consequence and incompatibility are related by the implicational explosion of inconsistent sets of sentences: inconsistent premise-sets imply every sentence. For both principled and empirical reasons, we deny that *any* of these global structural constraints should be taken to characterize reason relations in general, as we understand our topic. So we cannot appeal to this venerable idea to demarcate reason relations. How else, then?

Our answer, in short, is that reason relations are the common topic of a specific constellation of *meta*vocabularies. These are *rational* metavocabularies, in that the reason relations among their sentences depend, in various ways characteristic of the kind of rational metavocabulary they are, on the reason relations of the base vocabularies for which they are metavocabularies. Vocabularies whose reason relations do not depend on the reason relations of any other vocabularies are not rational metavocabularies in this sense—though to be vocabularies at all they must consist of sentences standing in reason relations. The ultimate goal is to understand the sentences of these base vocabularies—not just the reason relations among

them, but also their *use* to make claims and the *meanings* they express concerning how things are with the subject-matter they make it possible to talk about. We aim to illuminate those important topics by considering the reason relations the sentences stand in to one another, that is, the vocabulary that comprises those sentences. And the strategy is to understand base vocabularies—their use, the meanings of their expressions, and, as our central focus, the *reason relations* among sentences and the *conceptual roles* the sentences play in virtue of standing in those reason relations—in terms of *meta*vocabularies that have the expressive power to specify those uses, meanings, reason relations, and conceptual roles.

This strategy for specifying this new topic of discursive reason relations, and the notion of a *vocabulary* defined in terms of it, is *metalinguistic functionalism*. We understand reason relations in terms of the relations that articulate a whole constellation of *rational metavocabularies*: vocabularies for specifying reason relations. Reason relations, we want to say, are the common topic of just this constellation of metavocabularies. Our selection of the metavocabularies to consider is principled, and the space in which their relations to one another show up has a clear structure.

They are not, however, all created equal. We will take a frankly inegalitarian attitude toward them. The base vocabularies it is most important to understand are those codifying the reason relations of autonomous discursive practices—language-games one could play though one played no other, discursive practices that could be engaged in by practitioners who engaged in no others that were not fragments of them. Paradigmatically, these are natural languages. Picking them out in terms of their practical autonomy is demarcating them in a *pragmatic* metalanguage. Any such autonomous discursive practices must confer on some performances the significance of claimings, sayings in the sense of assertings. We think that means that they must also contain *reasoning* practices, in the form of claimings that serve as challenging other claimings by offering reasons against them, and other claimings that serve as defending them, by offering reasons for them. The relations between *reasoning*, as something done, and the reason *relations* of incompatibility (structuring reasons against) and implication (structuring reasons for) that are our proximal targets are complex. But the case for calling them *reason* relations ultimately turns on their relation to *practices* of rationally challenging and defending claims. Pragmatic metavocabularies, which specify reason relations in terms of the use of sentences to make claims and defend and challenge them, accordingly occupy a privileged place among rational metavocabularies. The substantive account they offer of what reason relations are is where we begin.

Representational semantic metavocabularies in the model-theoretic tradition offer a very different kind of substantive account of consequence

and incompatibility. This *Introduction* began with the thought that one might usefully illuminate the crucial relations between use and meaning (each the subject of an important philosophical tradition) by better understanding the relations between what is expressed by pragmatic and semantic metalanguages. Constructing, as we do, an isomorphism between suitably chosen versions of these kinds of metavocabulary shows that they indeed have a common topic. The content of identity claims, Frege taught us, is that two different senses pick out a common referent. Isomorphisms are equivalence relations, and the process of treating an equivalence as an identity is the essence of abstraction. The above-mentioned isomorphism accordingly shows that bilateral normative pragmatic metavocabularies and truth-maker semantic metavocabularies indeed have a common abstract topic: relations of consequence and incompatibility. That is what we mean by "reason relations."

Identifying this common target for accounts couched in these two principal kinds of extrinsic-explanatory metavocabulary brings into relief the possibility of another genus of rational metavocabularies addressing the same subject-matter. Intrinsic-explicative rational metavocabularies confine themselves to the conceptual resources supplied by the base vocabularies they address, and seek merely to express, rather than to explain, their reason relations. Logical vocabularies make explicit reason relations of implication and incompatibility in any arbitrary base vocabulary by extending it with conditional and negated sentences formed ultimately from the lexicon of that base vocabulary and then computing the reason relations of the extended vocabulary from those of the base. The implication-space rational semantic metavocabulary provides the expressive power to construct and manipulate the conceptual roles sentences play in virtue of standing to one another in the same reason relations articulated by logical vocabulary. These two principal kinds of intrinsic-explicative rational metavocabulary are related to each other not by the sort of isomorphism that relates the (right kind of) semantic and pragmatic extrinsic-explanatory rational metavocabularies. They are related to one another rather by the demonstrated soundness and completeness of the implication-space semantics for the most flexible and expressively powerful logic.

The internal structure of this constellation of two families comprising four kinds of rational metavocabulary is completed by showing how to construct an implication-space model corresponding to any truth-maker model of an arbitrary base vocabulary, by a mapping that preserves reason relations and conceptual roles. Doing that concludes the argument that it is the same thing that is explicated (each in its own way) by logical and implication-space metavocabularies, and that is explained (each in its own

way) by semantic and pragmatic metavocabularies: the reason relations and conceptual roles of sentences.

The illumination we believe is offered by telling the story we do here is of two kinds. The more general is the precise metalinguistic functionalist identification of a philosophical subject-matter—namely reason relations and the conceptual contents they articulate—as the common topic of metavocabularies of these four kinds, related to one another as we have shown they can be. The more specific comprises the particular versions of each of these kinds of metavocabulary that we elaborate: the more fine-grained bilateral deontic normative pragmatic metavocabulary, the specific consequence relation and way of relaxing global structural constraints in truth-maker semantic metavocabularies, the nonmonotonic multisuccedent logic that is provably expressively complete when extending even radically open-structured base vocabularies, and the implication-space semantic metavocabulary for conceptual roles. The details of those four particular ways of working out general conceptions of the expressive tasks characteristic of the different kinds of metavocabulary are essential to showing that they can stand in the relations to one another that reveal and delineate rational forms as their common topic.

The illumination an account offers of the phenomena it addresses crucially depends on its achieving the right level of abstraction. It must be general enough to provide a satisfying unification of otherwise disparate phenomena while specific enough to engage convincingly with the distinguishing details of the items it brings together. The metalinguistic functionalist picture that is elaborated in the body of this book is our candidate for achieving such a balance. The phrase "level of abstraction" is a loose way of talking about a structured set of isomorphisms. That is exactly how we have articulated the functional system of rational metavocabularies.

Every kind of rational metavocabulary is a form of rational self-consciousness. They are characteristic ways of talking and thinking about the rational structure of talking and thinking. For each is a distinctive way of expressing the reason relations that articulate the conceptual contents sentences acquire by being used as they are and standing to worldly states in the representational relations they do. Laying out the fundamental relations among these different kinds of rational metavocabulary is accordingly unpacking the internal structure of rational self-consciousness as such.

Notes

1 "Thought" in the sense of "thinkable," not in the sense of episodic "thinkings" of it (Frege, 1956).

2 A closely related result was first proven in (Hlobil, 2022a).

1 A Pragmatic Metavocabulary

Robert B. Brandom

1.1 Semantic and Pragmatic Construals of a Fundamental Bipolarity

The Fregean semantic tradition treats the distinction between *truth* and *falsity* as conceptually fundamental. This is not to say that these notions need be taken to be conceptually primitive, in the sense of not being further analyzable—though that was Frege's own view. But it is characteristic of this tradition to understand specifically *semantic* significance in terms of the difference something makes, its relevance to, possession of one or the other of these two "truth-values." The basic form of semantic meaningfulness is truth evaluability.

The two truth-values are not merely different (as **square** and **red** are), but exclusively different (as **square** and **circular** are). In the central, defining cases, possessing the one truth-value rules out possessing the other. It is impossible for anything to be both *true* and *false*. Further, the two truth-values are thought of as *opposites*. Appealing to an older tradition, we could say that "true" and "false" express Aristotelian contradictories rather than merely contraries. However, at least for the Fregeans, the opposition between truth and falsity is conceptually more basic than, and presupposed by, the distinction between contradictories and contraries. Much, much later in the order of explanation, sophisticated semantic theories might find reasons to relax these structural constraints for special cases, and countenance truth-value gluts and gaps: truth-evaluable items of which one says that they are either both true and false or neither true nor false. But the intelligibility of such late-coming subtleties is understood to be rooted in and parasitic on the more basic cases where what is truth evaluable takes one and only one of the two values. The home language-game of "truth-value" talk is built on and builds in both the distinction and the modally robust opposition between truth and falsity.

This familiar and influential philosophical tradition that has *semantic bivalence*, in a broad sense, as one of its guiding ideas is further characterized by its commitment to an order of explanation that seeks

to understand *meaning* ultimately in terms of truth-values. The purest, paradigmatic application, and a strong confirmation of the power of this approach, was provided early on by the clear, elegant, arbitrarily recursively extensible, bivalent truth-functional semantics it turned out to make possible for classical (Boolean) sentential *logical* vocabulary. Ever since, the first criterion of adequacy for a candidate semantics has been its capacity to provide a sound and complete specification of specifically logical consequence relations.

This bivalent semantic structure has an analogue on the side of pragmatics: that is, in theories of the *use* of linguistic locutions, rather than of the *meanings* or *contents* those locutions express. Here the starting point is *acts*, things speakers and thinkers can *do*, and the practical attitudes those acts express—in a different, pragmatic rather than semantic, sense of "express." The fundamental acts are *affirmation* and *denial*, expressing practical attitudes of *acceptance* and *rejection* (whether overtly and publicly, or covertly and quietly).

The acts expressing attitudes of acceptance and rejection can be as simple as saying "yes" or "no" in response to a claim or a question. These forms of affirmation and denial underscore that an essential element of the attitudes of acceptance and rejection is their opposite valence, their exclusion of one another. In the central, defining cases, one may say either "yes" or "no," but not both. The intelligibility of the cases where we do want to affirm and deny, for instance by saying "yes and no," is parasitic on the more basic cases where the opposition holds. Affirmations and denials count as *communicating information* only when the attitudes they express are treated as ruling one another out.

The sense of exclusion that applies to practical acts and attitudes is normative. One *ought* not to perform speech acts that would express acceptance and rejection of the very same claimable—even if one *could* do so. The deontic flavor of the modality of exclusion relating practical attitudes of acceptance and rejection on the side of pragmatics contrasts with the alethic flavor of the modality of the exclusion relating truth and falsity on the side of semantics. That exclusion is understood in terms of the *impossibility* (in the central, defining cases) of the very same truth-evaluable item being both true and false.

The semantic and pragmatic oppositions are linked by the possibility of understanding the practical attitude of accepting as *taking* (to be) *true*, and the practical attitude of rejecting as *taking* (to be) *false*. Such an identification is made possible by the fact that tokenings of declarative sentences can be understood as expressing both what is semantically truth-evaluable and what one can pragmatically accept or reject (affirm or deny, agree or disagree with, assent to or dissent from, say "Yea" or "Nay" to). Indeed, declarative sentences can be functionally specified just as what

can play both these roles. So construed, sentential lexical items are the syntactic glue that binds the semantic and pragmatic aspects of discursive practice. Together, these three form what we might call the "apophantic triad" of declarative sentences (syntactic characterization) expressing what is both evaluable as true/false (semantic characterization) and toward which one can adopt practical attitudes of acceptance/rejection (pragmatic characterization).

That those attitudes can be thought of as truth assessments or evaluations—that acceptance can be understood as taking true and rejection as taking false what is expressed by declarative sentences—opens up the possibility in principle of two orders of explanation. The semantics-to-pragmatics order of explanation would begin with a prior grasp of the concepts of truth and falsity and elaborate from that an account of what one needed to *do* to count thereby as practically *taking* or *treating* what is expressed by a declarative sentence *as* true or false, thereby accepting or rejecting it. (A potential candidate on the side of acceptance might include: using it as a premise for inferences, including practical ones issuing in intentional doings.) The pragmatics-to-semantics order of explanation would begin with a prior grasp of the concepts of acceptance and rejection and elaborate from that an account of what property one is thereby practically taking or treating what is expressed by a declarative sentence *as*. (A potential candidate on the side of truth might include: whatever is preserved by good inferences.) If neither of these explanatory strategies seems feasible (or attractive), one might treat the semantic and pragmatic dimensions of discursive practice as explanatorily coeval, and explain why they line up structurally as they do: true is to false semantically as acceptance is to rejection pragmatically.

How one understands these relations is a matter of some philosophical significance. For the opposition between truth and falsehood can be understood as a matter of how things are in the objective world one is talking or thinking *about*: how it is with the represent*ed*s one is representing in thought and talk. In contrast, the opposition between acceptance and rejection applies to how one represents things to be, that is, to the subjective side of the activity of represent*ing*. Insofar as we are interested in how these are related to one another, we should be interested in the structural *apophantic bipolarity* that shows up in the two interrelated species of the *semantic bivalence* of true and false and the *pragmatic dimorphism* of acceptance and rejection. Both, we have observed, are related by a modally robust kind of incompatibility. The semantic distinction of truth-values is articulated by the alethic *impossibility* of the same truth-evaluable being both true and false (in the same situation and at the same time). The pragmatic distinction of practical doxastic attitudes is articulated by the deontic *impropriety* of

the same claimable being both accepted and rejected (by the same subject at the same time). In either case, the *opposition* semantically between true/false and pragmatically between acceptance/rejection is crucial to one's understanding of thought and talk. The *division* into ways the world could be or opposing attitudes one could take to it is fundamental to either approach. Even if we want to allow qualifications around the edges, it is an essential element of the basic picture that in central cases the values or attitudes *exclude* one another.

The story just sketched is intended as a persuasive redescription of familiar ideas. The particular shape it takes serves two purposes within our larger story. To begin with, it provides the background for the pragmatics-first order of explanation elaborated in the rest of this chapter. More deeply, it sets the stage for what we take to be one of the most philosophically resonant and illuminating technical results that we have to report. This is the construction of a detailed isomorphism between a version of Kit Fine's model-theoretic hyperintensional truth-maker semantics and the version elaborated here of a bilateral normative pragmatics inspired by the work of Greg Restall and David Ripley. It is the topic of Chapter Four, at the heart of this book. The new way of looking at the relations between semantics and pragmatics (in our usage, the study of the meanings of linguistic expressions and the study of their use) that it offers sheds new light on the relations between representational conceptual content, the worldly states represented by such contents, and the activities that amount to representing them. It also makes visible a profound connection between what is expressed by the use of alethic modal vocabulary in (the right kind of) semantic theories and what is expressed by the use of deontic normative vocabulary in (the right kind of) pragmatic theories.

It might be noticed that the story told so far about semantic values of truth and falsity and practical attitudes of acceptance and rejection has not mentioned *negation*—and this might seem a culpable omission. After all, isn't the falsity of a claimable-believable-thinkable (what is expressed by a declarative sentence) just the truth of its negation? Isn't rejection or denial of it just acceptance or affirmation of its negation? Indeed it is. And large swathes of the philosophical tradition have succumbed to the temptation to bring negation into the story at the ground-level. But we will be rewarded by exploring a direction of explanation that understands negation in terms of the relations between the semantic distinction between truth and falsity and the pragmatic distinction between acceptance and rejection, rather than the other way around. That is the essence of the *logical expressivist* approach to demarcating logical vocabulary and concepts, which will be explained and motivated beginning in this chapter and the next, and then implemented in detail in Chapter Three. In accordance with that strategy, we proceed by considering some important features of the structure of

reasons and reasoning that can be made visible in advance of introducing logical vocabulary and concepts.

1.2 Reasons and Reason Relations: Symmetries and Asymmetries

Understanding the pragmatic attitudes (acceptance/rejection) and the semantic values (true/false) as related by the principle that acceptance is taking-true and rejection is taking-false is made attractive by the fact that it seems that *what is* accepted or rejected and *what is* true or false are the *same kinds* of things. What one can *take* true is what can *be* true. Very austere theories might have only *sentences* as what take truth-values and are accepted or rejected. (Quine is an example of such principled austerity.) When introducing the point above, though, it seemed natural to talk about what is true or false and what can be accepted or rejected rather in terms of what is *expressed* by declarative sentences. The thought is that the sentences themselves are candidates for semantic evaluation and pragmatic endorsement only at one remove, in virtue of what they express. The sentences are understood as expressive vehicles for what is in the primary sense truth-evaluable and affirmable or deniable: the object of pragmatic attitudes.

Why introduce this additional complicating factor? Because different sentences get truth-valued together, and accepted or rejected together, across a variety of situations. I say "I am confused on this point," and you say "You are confused on this point," or "Brandom is confused on this point." At noon I say "It is now noon," and 10 minutes later I say "It was noon then (when I last said something)." Ludwig wrote "Wovon man nicht sprechen kann, darüber muss man schweigen," and I say "About that of which one cannot speak, one must remain silent." When properly understood, the sentences (tokenings in the first two examples and types in the third) in these pairs stand or fall together, semantically and pragmatically. In this sense, sentences stand in *samesaying* relations to one another. This observation is the origin of the idea that what is evaluated semantically and pragmatically is what is *shared* by the sentences that belong to the same samesaying equivalence class. That is what invites us to think of those sentences as all *expressing* a common *content*. That expressive relation to a common content is understood as what determines that declarative sentence tokenings, or even types, *are* samesayings—that they stand in the samesaying relation to one another.

There is another sort of covariation of truth-evaluations implicit in grouping sentences into samesaying equivalence classes accordingly as they get the same semantic truth-evaluation in the same situations. For we can look at grouping those situations or circumstances into equivalence classes accordingly as they yield the *same* truth-evaluation for all the sentences of

some samesaying class. This is the origin of the idea of truth *conditions*, or of truth-makers (dually, falsifiers) associated with the contents expressed by the sentences of that samesaying class. This assimilation is not in the first instance of sentences, but of worldly states or situations in virtue of which (samesaying equivalence classes of) sentences are true or false. The work of the vague Scholastic metaphor "in virtue of which" in this claim is sometimes done by talk of what *makes* (samesaying equivalence classes of) sentences true or false. The more concrete metaphor immediately requires a qualifying commentary to the effect that the sense of "making" appealed to is not in general the *causal* sense of a temporal process that yields a result or effects a change. Moving in the other direction, the metaphor can be emptied out to the abstract form of what "determines" truth-values that invokes nothing more than a function in the mathematical sense, whose values (for some worldly arguments) are truth-values. In between is assertion of a modally robust covariation, distinguished by its support of subjunctive conditionals of the form "If the worldly situation *were* (or had been) thus-and-so, the (samesaying equivalence class of) sentences *would be* (or would have been) true (respectively: false)."

We might use the term "truth conditions" generically for the equivalence classes of worldly items playing this semantic explanatory role with respect to the samesaying equivalence classes of declarative sentences (tokenings and types). Some conceptions of truth conditions appeal to how it is with the worldly items the sentences (or the contents they express) are understood to be *about* (represent, denote, or refer to). On other conceptions, what makes sentences true are thought of as *facts*. (This is a non-Fregean sense of "fact," since in his usage a fact just is a thought—thinkable—that is true, not what *makes* it true or "corresponds to it," in any sense.) One form broadly truth conditional approaches to semantics can take is to construe the truth conditions of sentences as *intensions*: functions from indices to truth-values, thought of as the extensions of sentences. The worldly items that serve as arguments for such intension functions, and hence are assimilated as truth-makers or falsifiers of sentences, can be complete possible worlds or more fragmentary states or situations.

Thus far the emphasis has been on the structural homologies between talk of paired truth-values true and false in a semantic metavocabulary and talk of paired practical attitudes of acceptance and rejection in a pragmatic metavocabulary. There are structurally similar reasons to look for something in addition to sentences on the side not of what *is* true/false or accepted/rejected, but what in some sense *makes* (what is expressed by) declarative sentences true or false and *makes* attitudes of acceptance or rejection appropriate or correct. There is a specifically semantic sense of correctness of acceptance/rejection that can just be read off of evaluations

of truth and falsity, via the principle that accepting is taking-true and rejecting is taking-false. Acceptance is *semantically* correct just in case what is taken-true is true, and rejection is semantically correct just in case what is taken-false is false. And this non-causal, non-processual sense of "makes correct" can be expressed nonmetaphorically by a set of subjunctive conditionals: if the claim *were* true (its truth-conditions *were* satisfied), then acceptance *would* be semantically correct.

But there is another sense of "correctness" that applies to discursive practices, which confers on some acts and attitudes the practical significance of accepting or rejecting claimables (truth-evaluables) expressed by (what thereby count syntactically as) declarative sentences. It concerns *justification*, rather than truth. This is a still-normative notion of correctness that is neither equivalent to the semantic one nor reducible to mere matter-of-factual identification of the psychological causal antecedents of actual adoption of practical attitudes of acceptance and rejection: what makes practitioners adopt attitudes of acceptance or rejection in the sense of *causing* them to adopt those attitudes, or the processes by which they do. It is a matter, rather, of the *reasons* interlocutors have for the attitudes of acceptance and rejection they adopt and express (in a pragmatic, rather than semantic sense) in their acts of affirmation and denial.

Appeals to the reasons to accept and reasons to reject various claimables that are available to discursive practitioners can play explanatory-explicative roles that are analogous in some important ways to the roles played by appeals to intensions and truth conditions. It is clear that there will be important structural disanalogies as well. For the reasons available to an interlocutor might justify neither acceptance nor rejection of some claimable, and one might have both some reasons to accept and some reasons to reject some (other) claimable.

We began by pointing to an analogy between the traditional semantic extensions of declarative sentences, the two basic truth-values, on the one hand, and the two fundamental pragmatic acts and attitudes articulating the use of such sentences, acceptance and rejection (practically *taking*-true and *taking*-false), on the other. The possibility is now being raised of extending that analogy from the extensional to the intensional, by looking at the *reasons* interlocutors can have to adopt practical attitudes of acceptance and rejection. The first step on the path to elaborating pragmatics by analogy to the way semantic intensions are elaborated from merely extensional truth-values is firmly to distinguish practices of reason*ing* from reason *relations*. This is to distinguish reasons *to* do something—adopt an attitude or perform an act, accept or reject—from reasons *for* and *against* claimables (what is expressed by declarative sentences in the sense of what can be accepted or rejected).

Gilbert Harman (1984) vividly illuminates the required distinction as part of his argument for the initially shocking claim that there are no rules of deductive reasoning. Paraphrased in the idiom used here, he argues that if there were, presumably a paradigm would be the rule that if you accept p and accept $p \rightarrow q$, then you have decisive, deductively good reasons to accept q. So in those circumstances, you ought to do so. But, he points out, that would be a terrible rule. You might have much better reasons to reject q than you have to accept p or $p \rightarrow q$. In that case, you should reject one of them. Acceptance of conditionals can be exploited in reasoning either by *modus ponens*, or, equally validly, contrapositively by *modus tollens*.

What deductive logic directly supplies is reason *relations*. They tell us that some claimables provide reasons for and against others. Those relations are indeed relevant to practices of reason*ing*, but only indirectly. They constrain but do not direct the drawing of conclusions, the adoption of some doxastic attitudes as justified by the adoption of others. The fact that p and $p \rightarrow q$ stand in the relation of *implication* to q tells us that one *ought not* accept p and $p \rightarrow q$ and reject q. (We will be normatively "out of bounds" if we do.) But it does not tell us what to *do* should we find ourselves with those attitudes—which one or more of them we should change. The fact that p and $\neg p$ stand in the relation of *incompatibility* (here, formal incompatibility: inconsistency) tells us that we ought not to accept both. But again, it does not tell us what to do in such a situation—which attitude we should give up.

Harman is concerned with the role of specifically *logical* reason relations. But claimables (acceptable/rejectables) expressed by sentences that contain no logical vocabulary also stand to one another in relations of implication and incompatibility. (Following Wilfrid Sellars, we may call these *material* reason relations.) "Plane figure A is square" implies "Plane figure A is polygonal." So one ought not to accept the first and reject the second. "Plane figure A is square" is incompatible with "Plane figure A is circular." So one ought not accept both. Nested kinds support implications: "Coda is a dog" implies "Coda is a mammal," which implies "Coda is a vertebrate," and so on. And contrary properties support incompatibilities: "Coda is a mammal" and "Coda is a reptile," "Monochromatic figure A is green," and "Monochromatic figure A is red," and so on.

Harman argues that from relations of implication and incompatibility we cannot directly read off which acceptances and rejections give us reasons to accept and reject other claimables. We might, however, hope to be able to go the other way around. This would be to read off relations of implication and incompatibility between claimables from how some attitudes provide reasons to adopt other attitudes, via the non-determinative normative constraints that the reason relations impose on rationally admissible changes of attitudes. Pursuing this order of explanation would take us from

practices of giving and assessing reasons to accept and reject the claimables expressed by declarative sentences to reason relations of implication and incompatibility that those claimables stand in to one another. The aspiration would then be to understand those claimable contents—what can be accepted or rejected in reasoning—in terms of the reason relations they stand in to one another.

One way to begin is to take it that

- If accepting *A* functions practically as a reason to accept *B*, then *A* provides a reason for *B*, and
- If accepting *A* functions practically as a reason to reject *B*, then *A* provides a reason against *B*.

As the next step, we can then think of *implication* relations as codifying reasons *for* and *incompatibility* relations as codifying reasons *against*.

The idea is to start with practices of reasoning, in the sense of practices of giving reasons that entitle one to attitudes of accepting and rejecting claims (manifested in acts of asserting and denying them). We can think of a dialogical situation, where those who accept or reject a claim can be challenged to defend that attitude, to *justify* it by offering *reasons* to accept or reject it. These practices of asking for and offering reasons *to* do something, to accept or reject a claim (claimable), must respect reason relations among claimables according to which some of them provide reasons *for* and reasons *against* others. These we understand as relations of material implication and incompatibility. What stand in *these* relations are not acts or attitudes, but claimable contents: what one can accept or reject (whether reasonably or not, depending on what reasons *to* adopt those attitudes one can offer).

We gestured earlier at theoretical reasons to think that reasoning practices must include the possibility of offering and assessing reasons to accept. This is to rule out the ultimate intelligibility of purely *skeptical* reasoning practices: practices that permit the adoption and justification only of attitudes of *rejection*. (The challenge of making sense of first-person practical reasoning, practical deliberation, was offered as a suggestive case in point.) In the present context, any such considerations provide reasons to think that reason *relations* must include *implications*, which codify reasons *for*. Reasons to think that anything intelligible as reasoning practices must also include the possibility of offering and assessing reasons to *reject* are not far to seek. These would deny the ultimate intelligibility of purely *dogmatic* reasoning practices: practices that permit the adoption and justification only of attitudes of *acceptance*. In the present context, any such considerations provide reasons to think that reason *relations* must include *incompatibilities*, which codify reasons *against*.

In a wonderful essay called "Why 'Not'?," Huw Price (1990) considers the practical deficiencies of what we are calling "purely dogmatic" reason-giving practices. He imagines "ideological positivists," who do not have a way of denying or rejecting a claim. They accordingly lack any practical acknowledgment of the *incompatibility* of two claims. (It will follow that in their logic they have no way of *negating* a claim—hence the issue of his title.) He illustrates why such practices wouldn't work with a nice dialogue:

> Me: 'Fred is in the kitchen.' (Sets off for kitchen.)
> You: 'Wait! Fred is in the garden.'
> Me: 'I see. But he is in the kitchen, so I'll go there.' (Sets off.)
> You: 'You lack understanding. The kitchen is Fred-free.'
> Me: 'Is it really? But Fred's in it, and that's the important thing.'
> (Leaves for kitchen.) (Price, 1990, 224)

Unless the claims we accept can *exclude* some other acceptances, they can't guide our actions. The essential conceptual starting-point of Shannon information theory—well upstream of the issue of how to *quantify* information—is the idea that if a message does not exclude some alternatives that were previously open, it conveys no information at all. We would learn *nothing* practically from finding out that there are reasons *for* someone to accept a claim—say, "Fred is in the garden,"—unless those same considerations can serve also as reasons *against* accepting some other claims—"Fred is in the kitchen,"—which accordingly count as *incompatible* with the original claim. That means that the very same claim that is a reason *for* one commitment must also be a reason *against* some others. Not only must it be possible to accept or to reject any claimable, in addition, adopting either of those attitudes towards a claimable must be able to serve *both* as a reason to *accept* some further claimables (a reason *for* them), and as a reason to *reject* some other claimables (a reason *against* them). What can be accepted or rejected must stand both in relations of implication and in relations of incompatibility. A discursive practice cannot be intelligible as articulated by one sort of reason relation unless it is intelligible as articulated by the other as well.

We have argued that what can be accepted must be capable of being rejected, and what can be rejected must be capable of being accepted, and that what can serve as a *reason to* accept some acceptable/rejectable must be capable of serving as a reason *to* reject other acceptables/rejectables, and *vice versa*. We have accepted Harman's argument that one must distinguish between norms governing conditional practical attitudes of acceptance/rejection and the reason *relations* that constrain, but do not determine those norms. As a result, we have argued, one must understand what can be accepted or rejected as standing in *both* sorts

of reason relations: *implications*, codifying reason-*for* relations, and *incompatibilities*, codifying reason-*against* relations. All of this is a way of implementing the strategy of appealing to practical attitudes of accepting and rejecting what is expressed by sentences in order to understand the acceptable/rejectable *contents* expressed by declarative sentences by looking first to *reasons* other attitudes provide to accept or reject. The connection permitting this transition is supplied by the principle that a reason *to* accept (adopt that attitude) is governed by a relation between reasons *for* the content or object of that attitude (what is accepted), and that a reason *to* reject (adopt that attitude) is governed by a relation between reasons *against* the content or object of that attitude (what is rejected).

In this way we move from the idea of practical attitudes providing reasons *to do* something (adopt other attitudes) to relations of implication and incompatibility (reasons for and against) relating what can now be understood as what those attitudes are attitudes towards. It is an explanatory advance from *pragmatics*, studying what one is doing in adopting discursive attitudes, to *semantics*, studying the contents of those attitudes. Those contents are now thought of as nodes in a network of relations of implication and incompatibility.

To say that a set Γ of acceptables/rejectables *implies* acceptable/rejectable A, we can write "$\Gamma \mid\!\sim A$." Use of the "snake turnstile" rather than the more familiar double turnstile $\models$ of semantic consequence or the single turnstile $\vdash$ of derivability reminds us that we are expressing *material* implications, not *logical* implications. (An implication is logically good in case it meets two conditions: i) it is materially good, and ii) its material goodness is robust under arbitrary uniform substitution of nonlogical vocabulary for nonlogical vocabulary.) To say that a set of acceptables/rejectables is *incompatible* with acceptable/rejectable A, we can write "$\Gamma\#A$." (For our purposes here it suffices to stick to the more familiar and intuitive single-succedent notation. We'll have something to say later about the multisuccedent analogues.)

So far the discussion of the attitudes of acceptance and rejection and (so) of reasons-for and reasons-against in the form of implications and incompatibilities has been reasonably even-handed. The picture has been symmetrical. There is a substantial structural asymmetry between the two kinds of reason relation however. Implication is not in general a symmetric relation. If, possibly in the context of Γ, A implies B, it does not at all follow that in the same context B implies A. $\Gamma, A \mid\!\sim B$ does not entail $\Gamma, B \mid\!\sim A$. "Pedro is a donkey" implies "Pedro is a mammal," but not *vice versa*. By contrast, incompatibility is *de jure* symmetric. $\Gamma, A\#B$ *does* entail $\Gamma, B\#A$. "Oscar is an octopus" is incompatible with (a decisive reason against, a dispositive reason to reject) "Oscar is a mammal," and "Oscar

is a mammal" is incompatible with (a decisive reason against, a dispositive reason to reject) "Oscar is an octopus."

We can ask: why is it that one reason relation is symmetric and the other not? Must it be so? What defect would a discursive practice have if it did not exhibit this structural asymmetry between the two kinds of reason relation? And if there must be such a structural asymmetry, is it necessary that it be reasons *for* (codified in implications) that are nonsymmetric and reasons *against* (codified in incompatibilities) that are symmetric? Is it so much as intelligible that for some discursive practices it should be the other way around?

In *Making It Explicit* Brandom (1994) explicated discursive practice in a normative pragmatic metavocabulary of *commitment* and *entitlement* to commitments—a framework whose rationale we shall return to in the next section, as part of an ongoing attempt properly to understand and articulate the "ought" in "thought." In our favored deontic terms, one understands what it is for *A* to be incompatible with *B* as commitment to *A* precluding entitlement to *B*. It is not that it is *impossible* to commit oneself to *B* by asserting it even though one is already committed to *A*. It is just that if one does, one has foregone the possibility of having the normative status of entitlement to *B*. Other interlocutors who are aware of the incompatibility will not treat one as justified in claiming *B*, as having an entitlement to it that might be inherited by others testimonially, for instance. There is much to recommend such an understanding. But it is at least not obvious on such an analysis why the fact that commitment to *A* precludes entitlement to *B* should entail that commitment to *B* precludes entitlement to *A*. It seems possible that these could come apart, that we should keep separate sets of books on whether commitment to *A* precludes entitlement to *B* and whether commitment to *B* precludes entitlement to *A*. But we don't find examples like this in the wild. Why not?

At the end of Section 1.3, we will offer our answer to this question. The argument is a pragmatic one, appealing to dialogic features of practices of making claims and giving and asking for reasons for them. That such an argument is available only from a pragmatic point of view—only considering the essential structure of this fundamental reason relation in the context of practices in which it articulates practical standards for the assessment of the normative significance of performances—is part of what recommends a pragmatics-first order of explanation. In the next section, we assemble some additional conceptual raw materials that are needed to mount that larger argument. But first, an objection from a contrary logic-first (what we'll call a "logicist") order of explanation should be considered.

It might seem that this issue about the symmetry of incompatibility is wholly an artifact of the pragmatics-first order of explanation pursued

here, and that it is spurious, so that its arising at all counts substantially against that way of proceeding. For the traditional semantic theorist who treats logical vocabulary as available at the explanatory ground-level has a quick answer to the question about the rationale of the symmetry of incompatibility. To say that p and q are incompatible is to say that $p \rightarrow \neg q$. That is truth-functionally equivalent to $\neg(p \wedge q)$ and to $\neg p \vee \neg q$, which are symmetric because conjunction and disjunction are. So from a truth-functional perspective, $p \rightarrow \neg q$ if and only if $q \rightarrow \neg p$. That is just the symmetry of incompatibility, which we can codify in the Sheffer stroke.

Doesn't that analysis answer the question asked and settle the issue raised about the symmetry of incompatibility? It certainly shows that the symmetry of incompatibility is built deeply into semantic bivalence. So *if* the truth-functional connectives are the right way to codify and express material incompatibility, then it is symmetric. But such connectives are the right way to represent this reason relation *only if* material incompatibility, like logical inconsistency, is symmetric. It does in fact seem to be, which is a point in favor of the bivalent logical semantic analysis. But that analysis presupposes the symmetry of incompatibility, rather than explaining it—at least from the pragmatic point of view. We could say "So much the worse for the pragmatics-first order of explanation." But if we dig deeper, we can find a satisfying normative pragmatic explanation for the necessity of the symmetry of incompatibility. The invisibility of the question from the bivalent semantic logical point of view then must count against that order of explanation.

One mark of that invisibility, or better, technique for achieving it, is notational. In Gentzen-style sequent calculus formulations of reason relations, there is no separate sign for the relation of incompatibility. Incompatibilities are notationally assimilated to implications. To say that Γ is incompatible with A (what we are expressing by "$\Gamma \# A$"), we write something like "$\Gamma, A \mid\!\sim$ " (notice the empty right side). This attributes a property, call it "incoherence" to everything on the left of the turnstile. This notation obviously builds in the symmetry of incompatibility, since "$\Gamma, A, B \mid\!\sim$ " says *both* that, in the context of Γ, B is incompatible with A *and* that, in the context of Γ, A is incompatible with B (and similarly for any element of Γ). It is a measure of the success of this notational sleight-of-hand in distracting from and hiding the commitment to the symmetry of incompatibility that Gentzen, who basically agrees with Tarski about the structural properties of logical consequence relations (even though he uses lists of sentences where Tarski uses sets), does not treat the symmetry of incompatibility as one of his structural principles of reason relations. He does not enumerate it alongside monotonicity and transitivity (his "Thinning" and "Cut"). This substantial structural commitment is invisible to him as a distinct issue. (He builds it into the structural

principle of "Permutation" as applied to the premises of the sequents with empty right-hand sides that he uses to encode incompatibilities.) A related issue, further downstream, is that even logicians prepared to relax the monotonicity requirement on consequence or implication relations typically do not—and some approaches, for instance in terms of defaults, cannot—accommodate also the possibility of nonmonotonic incompatibility relations.

1.3 Bilateral Pragmatic Renderings of Reason Relations

We began by contrasting a *semantic* order of explanation, whose key conception is a distinction between two truth-values, true and false, with a *pragmatic* order of explanation, whose key conception is a distinction between two practical attitudes, acceptance and rejection (corresponding to acts of assertion and denial). In both cases there is a need to postulate something beyond declarative sentences, something that is in some sense *expressed by* such sentences, that is what is in the first instance true or false, accepted or rejected: the bearers of the truth-values and the objects of the doxastic attitudes. A traditional semantic approach is to understand propositional contents as intensions: functions from a set of arguments to truth-values. We have been introducing an alternate approach, within the pragmatic order of explanation. It looks instead to understand the propositional contents expressed by declarative sentences in terms of *reasons* interlocutors can have to adopt the basic attitudes of doxastic acceptance and rejection.

The idea is that the practices within which acts and dispositions to act acquire the practical significance of doxastic acceptance and rejection are essentially, and not just accidentally, practices in which it is appropriate to ask for *reasons entitling* one to or *justifying* one in accepting or rejecting. We may use the term "claimable" as shorthand for "acceptable or rejectable" and (so) "assertible or deniable"—taking it that it is the *same* kind of thing that can be accepted (taken-true) that can be rejected (taken-false) and *vice versa*, just as it is the same kind of thing that can be true that can be false. (That it is the same kind of thing that can be both true or false, on the semantic side, and accepted or rejected, on the pragmatic side, is a further, potentially more controversial commitment.) Then the thought is that practices of accepting and rejecting, asserting and denying, are always also practices that include defending and challenging claimables, by offering reasons to accept or to reject them. Of course we can think of species of acceptance and rejection that are not like this. Voters might express their acceptance or rejection of a proposal by ticking one or the other of two boxes, with the question of their *reasons* for doing so not being at issue (at any rate at for this purely voluntarist phase of

the political process). But it seems clear that in order to understand *what* those voters are accepting/rejecting, we need to look beyond this aspect of their practices. The suggestion is that for this purpose it might be necessary (and perhaps even sufficient) to look to what would count as reasons to accept or reject the proposals they are voting on (and perhaps, what further acceptances and rejections the attitudes expressed by voting one way or another would provide reasons for).

We invoked an argument made by Gil Harman to point out the need to distinguish norm-governed practices of reason*ing*—*doing* something, taking up stances or positions, making moves to and from them, and giving reasons for and against them—from reason *relations* among the claimables that define those positions (nodes in the network of relations) and license those moves. The practices consist of making claims (asserting and denying claimables), and defending and challenging them by offering reasons to accept or reject them. Offering such a reason is just accepting or rejecting some other claimables (acceptables/rejectables), which stand in the right reason-relations to the one being defended or challenged. Those reason-relations are *implication* and *incompatibility*, relations in which some claimables serve as reasons for or against other claimables, respectively. The connection between offering reasons *to* do something, accept or reject, and the relations of being a reason for or a reason against a claimable consists, to begin with, in reasons *to* accept being reasons *for* the claimable accepted (implying it) and reasons *to* reject being reasons *against* the claimable being rejected (being incompatible with it). The culmination of this pragmatic order of explanation would be an account of claimables—the propositional contents that are expressed by declarative sentences and can be accepted or rejected—in terms of the role the sentences that express them play in reason relations of implication and incompatibility. What is envisaged is a pragmatic route to a kind of semantics, in the sense of an account of propositional contents (claimables).

Before considering that final step, in order properly to understand relations of implication and incompatibility (what is expressed by the snake and hash turnstiles in metalinguistic statements of the form "$\Gamma \mid\!\sim A$" and "$\Gamma \# A$") we must look more closely at the reasoning practices that they codify. For in the context of the pragmatic order of explanation being considered, the only grip we have on these relations is the role they play in practices of defending and challenging claims, by giving reasons to accept or reject them. Here there are two main points that we would like to argue for. The first point is that to be intelligible as practices of reasoning, in the sense of accepting and rejecting claimables and defending and challenging those stances with reasons for and against them, the participants in such practices must be understood as keeping track of two different normative statuses: the kind of *commitment* one undertakes or

acknowledges in accepting or rejecting a claimable by asserting or denying a sentence expressing it, and the sort of *entitlement* to that status or practical attitude that is at issue when *reasons* are offered for or against it. The second point is that there is an important dimension along which these two flavors of normative status have quite different structures. The basis on which commitments are attributed is atomistic, while the basis on which entitlements are assessed is holistic.

As to the first point, we can begin with the observation that accepting or rejecting a claimable, paradigmatically by asserting or denying it, is taking a stand on it, adopting a stance towards it. It is committing oneself with respect to it, in the way one would by saying "Yea" or "Nay" to it in response to a suitable yes/no question. On the side of uptake, what some other practitioner needs to be able practically to discriminate in order to count as understanding the speech act is *that* the speaker has committed herself (performed a committive act, expressed a doxastic attitude), *how* she has committed herself (which kind of attitude she has adopted and expressed: acceptance or rejection, a positive or a negative commitment), and *to what* she has committed herself (toward which claimable she has adopted a doxastic attitude by asserting or denying the declarative sentence she uttered).

What difference does it make whether an interlocutor can offer reasons to accept what he has accepted or to reject what he has rejected? The doxastic *commitments* involved, the stances taken up, the attitudes adopted, are the same either way. But it is also an integral feature of doxastic commitments that one's *entitlement* to those (perhaps carelessly undertaken) commitments is always potentially at issue. For in taking up a doxastic stance one renders oneself liable to demands for justification, for exhibition of reasons to accept or reject the claim one has accepted or rejected. One's *liberty* to commit oneself, to adopt that attitude and acquire that status, is not an enduring *license* to do so. Reasons matter because other practitioners must distinguish between the acceptances and rejections the speaker in question is *entitled* to, in virtue of having *reasons* to adopt those attitudes, and those the speaker is *not* entitled to, because the speaker is unable to defend those commitments by offering reasons when suitably challenged to do so. It follows that for each interlocutor there must be not only a difference between the attitudes (commitments) he has adopted and those he has not, but also, within those he has adopted, between those he is entitled to or justified in, has rational credentials for, and those that are *mere* commitments, bare of such accompanying entitlements. In *Making It Explicit* Brandom argues that what turns practically on one's entitlement or justification is the testimonial authority of one's act: its capacity to license others to adopt a corresponding attitude. The essential point is that in addition to the *committive* dimension of assertional practice, there is the

critical dimension: the aspect of the practice in which the rational propriety of those commitments, their justificatory status, is assessed. Apart from this critical dimension, the notion of *reasons* gets no grip. It gets its grip from those keeping deontic score on their fellow discursive practitioners, who treat a failure to satisfy the justificatory responsibility implicit in undertaking a doxastic commitment as undercutting the interpersonal authority such a commitment otherwise could exercise.

The second structural observation is that entitlements are fragile in a way that commitments are not. Our picture of discursive practice understands assertion as having a default-and-challenge structure. Assertions can come with default entitlements. But those default entitlements can be lost when the assertion is challenged by offering reasons against the claim. And then, if its authority is to be regained, that entitlement must be re-acquired or vindicated by defenses offering reasons for the challenged claim. Entitlements are vulnerable to being undercut by incompatible collateral commitments. The basic phenomenon here is twofold. It is not impossible for someone to be committed both to accept and to reject the same claimable. But, one cannot then count as *entitled* to those contrary commitments. For each commitment provides a decisive reason against the other. The contrary commitments might have arisen through affirmation and denial of the same sentence—or, more commonly, when one is a (possibly unacknowledged) consequence of other attitudes the subject has self-consciously adopted. The mutual repulsion between the incompatible commitments that is implicit in attitudes of acceptance and rejection takes place at the normative level of rational entitlements to those commitments.

The origin and paradigm of the incompatibility of commitments undercutting their entitlements is the normative collision that occurs when one accepts and rejects the same claimable. But the phenomenon is not limited to that original case where contrary attitudes are adopted towards one and the same claimable. One treats the contents of two claimables as *incompatible* just by taking it that commitment to one precludes entitlement to the other. One forfeits entitlement to one's commitments if one both affirms and denies (accepts and rejects) that the plane figure is a circle. But one incurs the same normative cost if one both accepts that it is a circle and accepts that it is a triangle. That is the practical normative significance of the fact that "*A* is circular" and "*A* is triangular" stand to one another in the reason relation of material incompatibility (Aristotelian contrariety): commitment to one precludes entitlement to the other.

The fact that claimables stand to one another in the reason relation of incompatibility—the fact that commitment to one can preclude entitlement to the other—means that there is a structural asymmetry between the normative statuses of commitment and entitlement, which articulate essential dimensions of the practice of giving and asking for

reasons, making claims and defending and challenging them. Knowing an interlocutor's attitude toward a claimable, whether they accept or reject it, is sufficient to settle their commitment with respect to it. But to assess their *entitlement* to that commitment we have to consult *all* their other commitments. It is not enough that they can cite collateral commitments that provide good reasons for the commitment in question. It is necessary in addition that they have not undertaken commitments that provide equally good reasons against it.

Following a suggestion by Ryan Simonelli, we can assemble these conceptual raw materials so as to characterize both incompatibility and implication in terms of some sets of commitments precluding entitlement to others.[1] In the normative pragmatic vocabulary put in play here, we can define the reason relations like this:

Implication (IMP):
$\Gamma \mid\!\sim A$ if and only if commitment to accept all of Γ precludes entitlement to *reject* A.

Incompatibility (INC):
$\Gamma \# A$ if and only if commitment to accept all of Γ precludes entitlement to *accept* A.

On this account, a reason *against* a *rejection* is an implication with that conclusion, since $\Gamma \mid\!\sim A$ says that commitment to all of Γ precludes entitlement to reject A. That is a reason *for* an *acceptance*. Dually, an incompatibility $\Gamma \# A$ exhibits its premises as providing both a reason *against* acceptance and (so) a reason *for* rejection.

Notice that these pragmatic readings respect Harman's point. They do not directly dictate what inferences one draws. They address merely the rational cotenability of various attitudes. We have suggested that there is a useful sense in which the semantic contents of the claimables (acceptables/rejectables) expressed by declarative sentences can be understood as consisting in the roles they play in reason relations of these two kinds. These definitions show how semantic contents in this sense can be understood in purely pragmatic terms of commitments and entitlements to accept and reject the claimables that stand to one another in relations of implication and incompatibility. We see here deep connections among the paired distinctions between acceptance/rejection, commitment/entitlement, and implication/incompatibility.

Although we have presented it for the more familiar single-succedent turnstile, this pragmatic definition of implication is recognizably a version of the bilateralist reading Greg Restall and David Ripley recommend in order to make sense of the multisuccedent relation of implication (Restall,

2005, 2009, 2013; Ripley, 2013). Their account has the immediate benefit of demystifying multiple conclusion implications, which many have found hard to parse and motivate. (Why is the comma conjunctive on the left of the turnstile, when combining premises, and disjunctive on the right, when combining conclusions?) They recommend that we understand what is expressed when we write "$\Gamma \mid\!\sim \Delta$" (to put it in the notation we are using here) for sets of sentences Γ and Δ, as the claim that the normative position of anyone who is committed to accept all the sentences in Γ and reject all the sentences in Δ is "out of bounds." This philosophically powerful pragmatic interpretation allows them to understand sequent calculi as consisting of rules that tell us that if some positions are out of bounds, then some others are also. It is then easy to see how logic, so construed, normatively *constrains* reasoners guided by the aim of remaining "in bounds," without pretending to dictate unique answers to the question of what they should *do*, how in particular they should alter their commitments, in particular situations. Simonelli translates their normative pragmatic idiom into the vocabulary of commitment and entitlement familiar from *Making It Explicit*, and then shows how to give a parallel explicit treatment of incompatibility (expressed in Gentzen's sequent calculi by sequents with empty right-hand sides) using the common structure of some commitments precluding entitlement to others.

Even if one acknowledges the felicity and fecundity of Restall and Ripley's bilateral interpretation of the multisuccedent turnstile, one might still find it puzzling that implication relations should be thought of in terms of preclusion of entitlement. The standard way of thinking about implication (implicit in Tarski's formal articulation of it and explicitly endorsed in *Making It Explicit*) seems rather to involve some commitments having other commitments as consequences. This is the analogue, for reason relations among *non*logical (logically atomic) sentences of deductive logical consequence. Is the account being put forward here committed to that idea just being wrong about implication relations? Is there really nothing to be made of that line of thought? Further, it seems that if Γ is incompatible with A then accepting it *commits* one to *reject* A. Here again, Simonelli shows us the way. He points out in effect that we can introduce a pragmatic sense of "implicit" according to the two principles:

Pragmatically Implicit Acceptance (PIA):
Any set of commitments that *precludes entitlement* to reject A thereby *implicitly* commits one to accept A.

Pragmatically Implicit Rejection (PIR):
Any set of commitments that *precludes entitlement* to accept A thereby *implicitly* commits one to *reject* A.

PIA together with IMP entails that if $\Gamma \mid\!\sim A$, then commitment to accept all of Γ *implicitly* commits one to accept A. (Here "implicit" derives from "implies.") PIR together with INC entails that if $\Gamma\#A$, then commitment to accept all of Γ *implicitly* commits one to reject A. In this way we can reconstruct what is right about the thought that implication is a matter of acceptance of some premises having commitment to acceptance of a conclusion as its consequence and incompatibility is a matter of acceptance of some premises having commitment to rejection of a conclusion as its consequence. The connection between those characterizations and the modified bilateralist normative pragmatic construal is provided by the implicit acceptance and implicit rejection principles PIA and PIR.

Why should we accept those principles? In what sense are commitments to accept and to reject implicit in preclusions of entitlement to reject and to accept, respectively? The claim of PIA is that if commitment to accept precludes entitlement to reject A, then that same commitment to accept implicitly commits one to accept A. One option, rejecting A, has been ruled out. One could remain agnostic, neither accepting nor rejecting. But that's not right. After all, one of the options has been ruled out. One cannot become entitled to reject A. The only option left standing, the only one available that one could potentially be entitled to is accepting A. By hypothesis, one has not yet explicitly done that. But that attitude of acceptance is implicit in the ruling out (as something one cannot be entitled to) of the only other option, in the sense that it is the only option left open. This is not the same as actually adopting the attitude, and that is what we mark by calling the commitment to accept "implicit," by contrast to the actual, explicit adoption of it. It seems clear both that this is an intelligible *pragmatic* sense of "implicit commitment to accept" and that calling it that is motivated by the rendering impermissible of the only other active option, rejection, and the consequent relative pointlessness of remaining uncommitted.[2]

With the distinction between commitments and entitlements on board, and the example of Restall's and Ripley's normative pragmatic bilateralism in mind, we are in a position to get a clearer view of the phenomenon that led Harman to distinguish reason relations such as implication from reasoning practices such as inferring in the sense of accepting or rejecting some conclusion on the basis of accepting a set of premises. Put in our terms, he points out in effect that the implication $\Gamma \mid\!\sim A$ need not entitle one who is committed and entitled to accept all of Γ to accept A, even though the implication implicitly commits him to it. For the interlocutor might have much better reason against A—since $\Delta\#A$—than Γ provides for it. If the interlocutor is also committed and entitled to accept all of Δ, then $\Gamma \mid\!\sim A$ precludes entitlement to reject A and $\Delta\#A$ precludes entitlement to accept it. Each explicitly precludes entitlement to the commitment that

the other implicitly requires. The reason relations determine that one may not draw both conclusions, for one is never entitled both to accept and to reject *A*. That position is "out of bounds." Commitments to which one otherwise would be entitled can have that entitlement undercut by collateral incompatible commitments. But the colliding implications and incompatibilities only preclude joint entitlements. They do not say what individual commitments one should accept or reject, what attitudes one should adopt or revise in response to that normative constraint, in the particular practical context of a determinate set of prior commitments. The important basic lesson Harman teaches about relations of implication and incompatibility merely constraining rather than dictating reasoning practices—how one ought to adopt further attitudes in the light of their rational relations to one's antecedent attitudes—does not at all turn on the specifically contrapositive forms of reasoning from implications that he uses to illustrate the point, and can easily arise in their absence.

Acknowledging the Harman point about the importance of clearly distinguishing between reason *relations* of implication and incompatibility, on the one hand, and reasoning *practices* of making inferential moves, on the other hand, is entirely compatible with insisting on an essential relation between them. For it is compatible with taking the central, paradigmatic case of the activity of inferring to be *explicitly* acknowledging commitments (to accept or reject) that are *implicit* in one's other commitments, in the literal sense of being *implied by* them. Adopting the practical attitude of accepting *A* as the expression of one's recognition that one's commitment to accept all of Γ precludes entitlement to reject *A* is inferring *A* from Γ—even though that is not the only form such recognition can take. We will see in the next chapter that there is an intimate connection between this pragmatic notion of a practical attitude (to accept or reject) being implicit in a set of other such practical attitudes and the semantic idea of practical doxastic attitudes not only having explicit conceptual content, but also having implicit conceptual content, determined by what they imply or exclude. Both senses of "implicit" are to be understood in terms of reason relations.

Formulating the issue in the normative pragmatic metavocabulary being recommended also sheds some new light on the puzzle about the symmetry of incompatibility raised earlier. It was mentioned then that thinking just in terms of commitment and entitlement does not evidently provide a reason why the fact that commitment to *A* precludes entitlement to *B* should entail that commitment to *B* precludes entitlement to *A*. It seems intelligible that one of these relations should hold without the other. Why shouldn't incompatibility be nonsymmetric, as implication is? If we look at the normative definitions IMP and INC we get the beginning of an answer to this question. On the side of implication, it is intelligible that

commitment to accept *A* should preclude entitlement to reject *B* without its being the case that commitment to accept *B* should preclude entitlement to reject *A*. That is exactly how it is with "Pedro is a donkey," and "Pedro is a mammal." If you accept that he is a donkey, you are not entitled to deny that he is a mammal. But you can legitimately accept that he is a mammal and deny that he is a donkey, since he might be a capybara. The implication "Pedro is a donkey" $\mid\!\sim$ "Pedro is a mammal" rules out the position in which one accepts that he is a donkey and denies that he is a mammal. It says nothing about the legitimacy of switching the doxastic valence of those attitudes.

By contrast, we are to read "Pedro is a donkey" # "Pedro is a capybara" as saying that acceptance of Pedro being a donkey rules out entitlement to accept Pedro being a capybara. By contrast to the case of implication, the valence of the commitments in the case of incompatibility is the same. The Simonelli reading of implication in terms of preclusion of entitlement is recognizably a version of the Restall-Ripley reading of one's position being normatively "out of bounds" if one both accepts the premises and rejects the conclusion. Applying and extending that model to the case of incompatibility understands INC as a version of what we would put in Restall-Ripley bilateralist terms as that the position in which one accepts both the premises and the conclusion is out of bounds. But since acceptance is involved in both cases, this is saying that accepting all of the elements of the *set* one gets by adding the conclusion to the premises is "out of bounds." That is just what Gentzen represents by a sequent with an empty right-hand side. That marks the *set* as incoherent. Incompatibilities read off of that incoherent set will be *de jure* symmetric.

This argument is not decisive. If $A\#B$, then commitment to accept *A* precludes entitlement to accept *B* *whether or not one is entitled to accept A*. It does follow that one cannot be entitled to accept both *A* and *B*. But it does not follow from that fact that mere commitment to *B* precludes entitlement to *A*. On the Simonelli reading, commitment to *A* precludes entitlement to *B*. (In the general—but still single-succedent—case, $\Gamma, A\#B$ if and only if $\Gamma, B\#A$.) The symmetry argument needs that one cannot be entitled to accept all of both the left-hand and the right-hand side of the # turnstile. The extension of the analogy with the Restall-Ripley bilateralist reading of the implication turnstile to a reading of the incompatibility turnstile is suggestive, but not coercive. It is probative, rather than dispositive.

We can use the considerations assembled here to complete the argument, making visible the nature of the pragmatic necessity that incompatibility be understood as a symmetric relation. It requires looking more closely at the dialogical pragmatic context in which those who give and ask for reasons entitling interlocutors to their commitments defend and challenge those commitments. The argument here, too, is due to Ryan Simonelli

(2023). The basic dialogic significance of showing that someone's position is normatively "out of bounds" or that they are not entitled to one of their commitments is presumably to oblige them to change those commitments. In the case we are addressing, this means withdrawing a commitment shown to be incompatible with others that interlocutor has undertaken. It turns out that nonsymmetric incompatibilities cannot serve this purpose. Suppose *A#B* but not *B#A*: commitment to accept *A* precludes entitlement to accept *B*, but commitment to accept *B* does not preclude entitlement to accept *A*. Now consider an objection to a speaker *S* who is already committed to accept *A* and who then asserts *B*. It is pointed out that *S*'s commitment to accept *A* is incompatible with commitment to accepting *B*, so that *S* cannot be entitled to that commitment to accept *B*. *S* might repair his normatively "out of bounds" situation in response to this objection by withdrawing the commitment to accept *B*. Or, *S* could hold onto the acceptance of *B* and withdraw commitment to accepting *A*. If *S* does either of these things, *S* will be back normatively "in bounds" as far as this incompatibility is concerned. But if *S* makes the repair by withdrawing commitment to accept *A*, in the case where *A* is incompatible with *B* but *B* is not incompatible with *A*, he can immediately reassert *A*, committing to accept it once again. Then it is not open to the objector to point to his acceptance of *B* as making this move illegitimate, a commitment to which *S* cannot be entitled. For that he is already committed to accept *B* is by hypothesis no objection to his acceptance of *A*. It does not preclude entitlement to that acceptance. The upshot is that nonsymmetric reason relations of incompatibility would be of no practical use in criticizing the commitments of any interlocutor. For any lack of entitlement they invoke can be repaired just by withdrawing the antecedent commitment and then endorsing it once again.

This argument is in some ways analogous to Dutch Book arguments in rational choice theories. It shows that those whose commitments are normatively criticized by invoking nonsymmetric incompatibility relations can immunize themselves from the effects of such criticism by the simple mechanism of withdrawing prior commitments shown to be incompatible with, and so to rule out entitlement to, subsequent commitments, and then reinstating those same commitments. It would be a mistake to think that this argument turns on contingent details of the dialogic rules for challenging an interlocutor's claims by offering reasons against them by exhibiting their incompatibility with other commitments—which, accordingly, preclude entitlement to those claims. One might be thinking of practices in which one was not permitted to withdraw earlier commitments, but "lost" the dialogical game if convicted of incompatible commitments. But the point goes deeper than that. Entitlement-precluding incompatibilities must be assessable with respect to the whole set of

an interlocutor's commitments. The question is whether any of them preclude entitlement to any of the rest. Nonsymmetric incompatibility relations would require distinguishing between two interlocutors who had exactly the same commitments, finding one "out of bounds" and the other "in bounds" normatively, depending on the order in which they had acquired those commitments. The requirement that incompatibilities be symmetric is the requirement that entitlements be assessable relative to the whole set of commitments whose entitlements are being assessed. That rules out one kind of what we might call "doxastic hysteresis," namely the path-dependence of entitlements consequent upon nonsymmetric incompatibility relations. We will see in the next chapter that there are other kinds of doxastic path-dependence, and that it is an important expressive criterion of adequacy for logical vocabularies that they are robust enough to help us reason in such situations.

1.4 Conclusion

In the very broad usage we give to the term here, "pragmatics" addresses discursive practice, the use of linguistic expressions. In the light of the circumscribed ambitions of this volume, we restrict ourselves to considering declarative sentences. Declarative sentences are picked out by their role in expressing practical attitudes of acceptance and rejection, by performances having the pragmatic significance of assertions and denials. One of our leading ideas is that that significance essentially depends upon discursive practices including assertions that stand to other assertions and denials as rational defenses of and challenges to them. Functionally characterizing one speech act as a *rational* defense or challenge is understanding it as providing a *reason for* or *reason against* the practical (we can now say "doxastic") attitude that is defended or challenged.

Our more proximate topic is what we call "reason relations" of implication and incompatibility, which generalize relations of deductive logical consequence and logical inconsistency to material non- or pre-logical reasons for and reasons against. As Harman reminds us, the connection between such relations and norms governing the making of claims and assessments of the success of rational challenges to and defenses of such claims is complex and indirect. The model we use to explicate part of the story is Restall and Ripley's bilateral normative construal of multisuccedent implication. According to that model, to say that one set of sentences Γ implies another set of sentences Δ is to say that the position of an interlocutor who asserts all of Γ and denies all of Δ is normatively out of bounds.

We refine and extend this account in two stages. First, we translate it into a two-sorted deontic normative vocabulary of commitments (to

accept and reject) and entitlements to those commitments, in place of the single-sorted normative vocabulary of "out-of-boundness" governing constellations of practical attitudes conceived in non-normative terms. Second, we adapt an idea of Ryan Simonelli to use this richer normative pragmatic metavocabulary to define reason relations of incompatibility as well as those of implication. These definitions in turn support a dialogic model of discursive practice, in which default entitlements to doxastic commitments ebb and flow through challenges and defenses that consist of offering reasons for and against contested claims. The sinews of this critical rational structure of core discursive practices are the reason relations of implication and incompatibility, which determine what commitments, expressible by sentences, provide reasons for and against (entitling one to accept or reject) what others. Our topic in the rest of the book is a constellation of formal metavocabularies for expressing and redescribing such reason relations, and the investigation of various illuminating relations among those rational metavocabularies.

Against the background of this critical linguistic rationalist picture of discursive practice, hints were offered about the prospects of a broadly inferentialist semantic alternative to truth-based semantics. Where the tradition advances from concern with bivalent truth *values* to a picture of meaning as truth (and falsity) *conditions*, the possibility was raised of advancing instead (or, as we will see, also) from a bilateral understanding of reason relations of implication and incompatibility to an understanding of meaning in terms of the roles sentences play in constellations of reason relations. In subsequent chapters we will make good on this suggestion in the form of a model-theoretic implication-space semantics for conceptual roles defined functionally by their reason relations of implication and incompatibility. We then explore both the relations between this sort of semantic metavocabulary and truth-maker semantics, on the one hand, and how both relate to the version of the bilateral normative pragmatic metavocabulary considered here, on the other.

So we have here both introduced the idea of understanding the central and characteristic *use* of declarative sentences in terms of *reasons* rather than in terms of *truth*, and at least gestured at the possibility of extending this approach to understanding in terms of reasons rather than truth also the *meanings* those sentences express by being used as they are in the making of claims. From a traditional point of view it is remarkable that, in addition to putting issues of truth to one side, the (admittedly abstract) discussion of reason relations involved no essential appeal to *logic* and logical vocabularies. It was acknowledged in passing that we think it is important not to restrict our conception of implication to formal logical consequence as deducibility or derivability and not to restrict our conception of incompatibility to logical inconsistency, though we did

not offer reasons in support of those commitments. We agree with the philosophical tradition not only that logic is intimately related to doxastic reasons, but also with the much stronger claim that one cannot in the end understand reason relations without understanding their relation to what is expressed by logical vocabulary. But we have a very different view about how that works and why that is so. That is the topic to which we turn next.

Notes

1 In presentations to the "Research on Logical Expressivism" (ROLE) working group.

2 The goal here is to articulate a motivation, and in service of that goal many considerations that would be relevant in other argumentative contexts are being suppressed. In particular, one might think that it is criterial of *paradoxical* sentences such as the Liar that subjects end up rationally committed *both* to accepting *and* to rejecting them, or that they are paradigms of sentences rational subjects should endeavor *neither* to accept *nor* to reject. Though we are not concerned to address semantic paradoxes in this work (save for suggesting in passing that they might be considered a species of a wider genus of expressive paradoxes), in Chapter Five we do consider three-valued logics LP (Logic of Paradox) and K3 (Strong Kleene) that have been appealed to in developing these ideas through gluts and gaps (for instance by Graham Priest and Saul Kripke, respectively).

2 Logical Expressivism and the Open Structure of Reasons

Robert B. Brandom

2.1 A Pragmatic Definition of Implication and Incompatibility

Chapter One introduced a pragmatic metavocabulary: a vocabulary for specifying discursive practices and the basic roles performances can play in such practices. Specifically *discursive* practices are distinguished as those in which some performances are practically accorded the pragmatic significance of *claimings*. The idea is that discursive practices can be picked out as practices of *explicitly expressing* doxastic practical attitudes. They are practices that permit one to *say that* things are thus-and-so. Declarative *sentences* are identified as linguistic expressions whose free-standing (unembedded) utterance has the default significance of claimings.

Claimings are understood as essentially *bivalent* speech acts. They take two complementary forms: *assertions* and *denials*. Assertions express practical attitudes of *acceptance* (of what is expressed by declarative sentences) and denials express practical attitudes of *rejection* (of what is expressed by declarative sentences). One might, but need not, in addition understand the attitude of acceptance as *taking-true* and rejection as *taking-false*. However on the pragmatics-first order of explanation adopted here, the pragmatic attitudes of acceptance and rejection, and the speech acts of assertion and denial that express them, are conceptually and explanatorily prior to the concepts of truth and falsity.

A further overarching commitment of the present approach is a kind of *pragmatic linguistic rationalism*. It affirms the *critical rational* character of basic discursive practices. This is the claim that it is essential to the pragmatic significance of claimings as such that they are subject to rational *challenge* by some other claimings, and in need of rational *defense* by yet different claimings. Challenging and defending can be thought of as analogous to illocutionary forces claimings can have. Challenging is making claims that serve as *reasons against* the challenged claim. Defending (justifying) is making claims that serve as *reasons for* the challenged claim. The idea is that performances can have the pragmatic significance of

claimings only in the context of practices of giving and asking for reasons. Claimings are what both can be given as reasons and for which reasons can be asked—what can both serve as and stand in need of reasons.

A *two-sorted deontic normative* pragmatic metavocabulary is apt for specifying the critical rational character of claimings. Claiming is understood as undertaking a distinctive kind of *commitment*. Commitments of this kind, *doxastic* commitments, come with a sort of default *entitlement*. The claimer is entitled to the doxastic commitment undertaken by claiming until and unless the claim is successfully challenged by claims expressing reasons against it. That entitlement can then be vindicated or regained by successfully defending it by claims expressing reasons for it.

In order to be intelligible as playing the role of claimables in critical rational discursive practices, what is expressed by declarative sentences must accordingly be understood as standing to what is expressed by other declarative sentences in two kinds of *reason relations*. Reasons *for* are claimables that *imply* the claimable they stand to as reasons for. Reasons *against* are claimables that are *incompatible with* the claimable they stand to as reasons against. The bivalence of doxastic attitudes of acceptance and rejection, manifested in the bivalence of speech acts of assertion and denial, together with the essential critical rational articulation of claimings as such together entail the *dyadic structure of reason relations*: their division into relations of *implication* and *incompatibility*. They may be thought of as relations of rational *inclusion* and *exclusion*. *Semantic inferentialism* is the thesis that the claimables expressed by the declarative sentences used to assert and deny are to be understood in terms of the roles those sentences play with respect to reason relations of implication and incompatibility. This is what we shall mean by "conceptual roles" and "implicational roles." Chapter Five will introduce what we claim is the intrinsic metavocabulary of conceptual roles: implication-space semantics.

The two-sorted deontic normative pragmatic metavocabulary is expressively sufficient to permit the *bilateral* definition of reason relations. A set of sentences Γ *implies* a sentence A ($\Gamma \mid\!\sim A$) if and only if *commitment* to *accept* all of Γ *precludes entitlement* to *reject* A. A set of sentences Γ is *incompatible* with a sentence A if and only if *commitment* to *accept* all of Γ *precludes entitlement* to *accept* A. If Γ *implies* A we can say that commitment to accept all of Γ *implicitly* commits one to *accept* A, equating preclusion from entitlement to reject with implicit commitment to accept. If Γ is *incompatible with* A we can say that commitment to accept all of Γ *implicitly* commits one to *reject* A. *Inferring* is *explicitly* acknowledging *implicit* commitments: actually accepting something that one's other commitments preclude one from entitlement to reject, or

actually rejecting something that one's other commitments preclude one from entitlement to accept.

2.2 Reasons and Logic: Logicism about Reasons

So far we have articulated a version of a *pragmatic discursive rationalist* view that puts reasons and reason relations of consequence and incompatibility at the center of an account of what distinguishes specifically *discursive* practices. We have also gestured at a *semantic inferentialist* order of explanation that seeks to understand the conceptual contents declarative sentences express in terms of the role those sentences play in such reason relations—though this will not be a central topic of this work until Chapter Five. This chapter and the next are devoted not to pragmatics or semantics, but rather to *logic*. They are animated by the conviction that an axial issue in the philosophy of logic is what might be called "the reasons question about logic." How should we understand the relations between logic and reasons?

There is a long tradition that takes logic to be a study of reasons. The reasons question about logic asks for further specification of the distinctive *kind* of inquiry into reasons logic is: its characteristic methods, goals, and presuppositions. Frege's 1879 *Begriffsschrift*—the founding document of modern logic—offers a resounding and sophisticated response to this challenge, of the expressivist sort that we will be recommending and pursuing. But his mature work, in the 1890's, shifts focus, from reasons to truth. Michael Dummett says of this shift:

> ... in this respect (and in this respect alone) Frege's new approach to logic was retrograde. He characterized logic by saying that, while all sciences have truth as their goal, in logic truth is not merely the goal, but the object of study. The traditional answer to the question what is the subject-matter of logic is, however, that it is, not truth, but inference, or, more properly, the relation of logical consequence. This was the received opinion all through the doldrums of logic, until the subject was revitalized by Frege; and it is, surely, the correct view.
>
> (Dummett, 1973, 432)

And again:

> It remains that the representation of logic as concerned with a characteristic of sentences, truth, rather than of transitions from sentences to sentences, had highly deleterious effects both in logic and in philosophy. In philosophy it led to a concentration on logical truth and its generalization, analytic truth, as the problematic notions, rather than

> on the notion of a statement's being a deductive consequence of other statements, and hence to solutions involving a distinction between two supposedly utterly different kinds of truth, analytic truth and contingent truth, which would have appeared preposterous and irrelevant if the central problem had from the start been taken to be that of the character of the relation of deductive consequence.
>
> (Dummett, 1973, 433)

Dummett is here acknowledging the centrality of the reasons question to the issue of what logic is. It is, he says, the study of a certain kind of *consequence* relation: "logical" consequence or "deductive" consequence. One way of sharpening the reasons question then is to ask how we should understand the relations between *logical* consequence or implication, and consequence or implication in *general*. Though he does not do so in these passages, presumably Dummett would acknowledge that logic also studies the notion of inconsistency, which is specifically *logical* incompatibility. And once again, we can ask how this special logical kind of incompatibility is related to the incompatibility of claimable conceptual contents in general.

The reasons question invites us to consider two contrary orders of explanation: understanding and explaining *reasons* in terms of *logical* reasons, and understanding *logical* reasons in terms of reasons more generally.[1] Frege pursued a logicist reconstruction of central mathematical concepts, explaining the reason relations they stand in to one another entirely in terms of specifically *logical* reason relations. His striking success (in spite of some attendant sophisticated and telling failures) was a testament to the expressive power of his new logic. It inspired a characteristic, axial project at the core of what was to become analytic philosophy: to reconstruct the reason relations that articulate conceptual contents in general using only *logically valid* relations of consequence and inconsistency. That is to do for important concepts outside mathematics what Frege had sought to do for central concepts within mathematics.

The master idea is that "*good* reason" just means "*logically* good reason." Where this is not obvious, it is the philosopher's task to unearth the hidden logical articulation of apparently nonlogical concepts, so as to be able to reconstruct, in purely logical terms, the relations of consequence and incompatibility they stand in to other concepts. The paradigm epitomizing this enterprise was Russell's 1905 essay "On Denoting." It discerns a logical deep structure underneath the definite descriptions of ordinary nonlogical descriptive vocabularies—expressions such as "the author of *Waverly*" and "the present King of France." Russell then shows how the consequences and incompatibilities of different kinds of sentences in which definite descriptions occur can be derived *logically*, as valid implications

and provable inconsistencies, from their suggested reformulations in logical terms.

The project launched by what Russell made of Frege is *logicism about reasons*, extended and generalized from its original form, which was addressed specifically to mathematical reasoning, to include the reason relations articulating *all* legitimate conceptual contents. Wittgenstein's *Tractatus* distills this logicism to its pure, crystalline form. In its subsequent empiricist form, logicism is the animating idea of the program of Carnap's Vienna Circle. Nonlogical content is provided to empirical concepts by their relations to perceptually occasioned noninferential observation reports, but the implications and exclusions of one another by those empirical concepts are to be shown to be matters exclusively of *logical* consequence and inconsistency, by suitable reconstructions or excavations of the logical form those concepts give to that raw nonlogical, nonconceptual, empirical content. According to this way of thinking about the relations between logic and reasons, logic is not just the *study* of reasons, it is the *science* of reasons. Logic sets the standards deciding what is a genuine reason for and against what, no matter the nonlogical subject matter being reasoned about.

Thought of this way, the logicist answer to the reasons question is recognizably a successor conception to what is arguably a founding idea of the Western philosophical tradition. The Greeks were puzzled and fascinated by the peculiar normative force of the better reason. How was this unique kind of specifically *rational* compellingness to be distinguished from or within the genus of utterances that merely reliably dispose their audiences to assent, without providing justifying reasons—the functional characterization of Sophists' techniques? Plato found his answer in the geometry that Euclid would go on to codify. For there, not acknowledging as compelling an argument in the form of a proof just shows that one did not understand it. In geometrical argumentation reasons display their force unhindered and unobscured. The philosophical enterprise is coeval with the aspiration to understand *all* reasons on this model.

The path from this mathematical model of reasons in geometrical form to a logicist model is long, and distinguished. While accepting Galileo's claim that the book of nature is written in the language of mathematics, Descartes generalized philosophically from his own analytic geometry and understood reasoning about the geometrical world of extension in terms of the manipulation of algebraic symbols. Spinoza hews to the geometrical model of reasons, subtitling his book on Descartes "*More Geometrico Demonstrata*," and his *Ethics* "*Ordine Geometrico Demonstrata*." Leibniz could envisage a future logic as a *calculus ratiocinator* only because he was looking beyond the syllogistic logic of the schools towards an ideal, specifically mathematical, logic in the form of an algebra. Nor was this line

of thought confined to rationalists. Hobbes expressed his endorsement of the mathematical model of reasoning in general in the slogan "Reasoning is but reckoning" (meaning arithmetic calculation). It was not until Frege's *Grundgesetze* and Russell and Whitehead's *Principia Mathematica* worked out in concrete detail projects for reconstructing substantial mathematical reasoning in terms of their new, expressively powerful logics that the general logicist program of understanding all theoretical reason relations as disguised logical reason relations could give shape and impetus to the emergent, ascendant school of analytic philosophy.

The generalization of logicism about specifically mathematical reasoning to logicism about reasons *tout court* was a bold and original philosophical thesis. Looking back over the results of more than a century's efforts in pursuing its promises and working this claim out in detail raises the suspicion that, as Hempel says of the Vienna Circle's attempts to formulate an empiricist criterion of meaningfulness, it

> … has come to play the role of the treasure hunt in the tale of the old winegrower who on his death-bed enjoins his sons to dig for a treasure hidden in the family vineyard. In untiring search, his sons turn over the soil and thus stimulate the growth of the vines: the rich harvest they reap proves to be the true and only treasure in the vineyard.
>
> (Hempel, 1963, 707)

We will see that the general logicist model of reasons faces deep difficulties of principle, rooted in the most general features of the structure of reason relations of consequence and incompatibility. Logicism about reasons entails that reason relations in general have (or can be reduced to logical relations that have) the topological *closure* structure characteristic of logical reason relations. If that *structural* logicist claim is not correct, then logicism about reasons cannot be right either.

2.3 Reasons and Logic: Rational Expressivism about Logic

Before turning to that structural argument against logic-first explanations of the relations between logic and reasons, it is worth considering for comparison a converse approach that pursues a complementary, reasons-first order of explanation. Logical *expressivism* is the view that logic is a *language* of reasons rather than a *science* of reasons. It is not a theory that determines or explains what is a reason for and against what. The distinctive task of logical vocabulary is not to *prove* that one set of claims implies or is incompatible with another, or to *derive* what is a reason for and against what, but rather to make it possible to *say* what is a reason for and against what. In this sense, logic is not a *canon*, but an *organon* of reason:

not a set of principles determining or explaining reason relations but a set of expressive tools for making reason relations explicit. Making implications and incompatibilities explicit is codifying them in the sentential form of claimables, that is, as the sort of things that can themselves be given as reasons and for which reasons can be asked. The point of adding logical vocabulary to a base vocabulary is to make it possible to talk about the reason relations of that base vocabulary—and so to reason about those reason relations, in the sense of making claims about them, and challenging and defending those claims.

Expressivism is the philosophical understanding of the characteristic, defining task of logic that Frege introduces in his seminal *Begriffsschrift* of 1879. To introduce his project he says:

> ... there are two ways in which the content of two judgments may differ; it may, or it may not, be the case that all inferences that can be drawn from the first judgment when combined with certain other ones can always also be drawn from the second when combined with the same other judgments. The two propositions 'the Greeks defeated the Persians at Plataea' and 'the Persians were defeated by the Greeks at Plataea' differ in the former way; even if a slight difference of sense is discernible, the agreement in sense is preponderant. Now I call that part of the content that is the same in both the conceptual content. Only this has significance for our symbolic language [*Begriffsschrift*] ... In my formalized language [BGS] ... only that part of judgments which affects the possible inferences is taken into consideration. Whatever is needed for a correct ['richtig', usually misleadingly translated as 'valid'] inference is fully expressed; what is not needed is ... not.
>
> (Frege, 1998, sec 3)

His logical notation is conceived as a tool for expressing the inferences sentences are involved in. And he introduces a special notion of content, *conceptual* content (*begrifflicher Inhalt*), which he identifies with the role sentences play in inferences (*Schlüsse*)—that is, their role in what we have been calling "reason relations." Expressing inferential relations is a way of making conceptual contents explicit. This is a semantic inferentialist thesis. It is why he calls his logic a "concept vocabulary," a way of expressing concepts in writing: *Begriffsschrift*. (His two-dimensional logical notation can *only* be written, not spoken.) Describing this project afterwards, Frege says:

> Right from the start I had in mind the expression of a content ... But the content is to be rendered more exactly than is done by verbal language ... Speech often only indicates by inessential marks or by imagery what

> a concept-script should spell out in full.
> (Boole's logical Calculus and the Concept-script, in Frege, 1979, 12-13)

This is logical expressivism: logic provides tools for characterizing reason relations, specifying what follows from what and what is incompatible with what.

It is important that as examples of sentences whose conceptual content (inferential roles) his logical vocabulary serves to explicate he offers *nonlogical* sentences, in the sense of sentences that contain no specifically logical vocabulary. The conceptual contents he aims to "spell out in full" to begin with are mathematical concepts concerning numbers. But he is clear that his expressive ambitions extend much further:

> It seems to me to be easier still to extend the domain of this formula language to include geometry. We would only have to add a few signs for the intuitive relations that occur there ... The transition to the pure theory of motion and then to mechanics and physics could follow at this point.
> (Frege, Begriffsschrift Preface, in van Heijenoort, 1967, 7)

He thinks that sentences with these nonlogical conceptual contents, for instance, those of physics, already stand in reason relations of implication and incompatibility to one another before logic comes on the scene. The task of logic is to "render those contents more exactly" by making explicit those nonlogical reason relations. Since he is a semantic inferentialist, understanding the conceptual contents of sentences in terms of their roles in reason relations, Frege thinks that by making it possible to specify reason relations clearly and precisely, logical vocabulary thereby enables perspicuous representation of the conceptual contents expressed by sentences formulated using nonlogical vocabulary.

According to such an expressivist view, the central, defining use of logical vocabulary is its use in conjunction with some nonlogical base vocabulary, whose reason relations it makes explicit. Frege explicitly mentions as examples of target base vocabularies, to begin with the vocabularies of arithmetic and geometry in mathematics, and the vocabularies of kinematics and dynamics in physics. Those base vocabularies are not just unstructured sets of lexical items. They are construed as coming with their own native reason relations of implication and incompatibility. These are what Sellars calls "material" inferential relations. He uses that term to emphasize that they are not relations of *logical* consequence and inconsistency. They articulate the conceptual contents of the *non*logical concepts that are deployed in the base vocabulary—as for instance the goodness of implications whose premises are of the form "*A* is to the East

of *B*," and whose conclusions are of the form "*B* is to the West of *A*," is part of what makes the nonlogical terms "East" and "West" mean what they mean, express the conceptual contents that they do. Newtonian physics determines what implications relate claims about the mass and acceleration of a body to claims about the forces acting on it. The expressivist idea is that logical vocabulary can then be added to such a base vocabulary, to codify its material reason relations in perspicuous form.

Of course, for the addition of logical vocabulary to a base vocabulary to make it possible to express the reason relations of the base vocabulary, the reason relations in which sentences of the logically extended vocabulary stand to one another must also be determined. Doing that is settling how the implications and incompatibilities of logically complex sentences formed by applying logical operators to sentences of the base vocabulary systematically inherit their reason relations from those of the base vocabulary. It turns out that that can be done, for instance in Gentzen-style sequent calculi, in a purely logical metavocabulary. That is, one can specify the reason relations that hold among sentences that contain no logical terms in such a way that it is settled thereby how the reason relations that hold between arbitrary sentences of *any* logically extended base vocabulary are determined by the material reason relations native to the base vocabulary. This is a remarkable fact, and a substantial achievement. (In Chapter Three we will look in some detail at how this works for our specific candidate for an expressively powerful logic.) It is this fact about the specifiability of purely logical reason relations, together with the fact that logical vocabulary so articulated can then be used to codify the reason relations of any base vocabulary whatsoever, that makes it tempting to conjecture that all material reason relations are made good by implications and incompatibilities that hold in virtue of logic alone—as the logicist about reasons does.

The logicist takes the home language-game of logical locutions—the sort of use that defines them as *logical* locutions—to be *pure* logic. On the side of the lexicon, pure logic is concerned with sentences that contain *only* logical vocabulary. In the sentential case, with which we are exclusively concerned in this work, this means sentential connectives, such as conditionals, negation, conjunction, and disjunction, together with schematic variables standing in for the sentences they are applied to, from which logically complex sentences are formed. And on the side of reason relations, the vocabulary of pure logic considers only relations of logical consequence and inconsistency that hold just in virtue of those logical connectives: for example the implication of A by $A \wedge B$ and the incompatibility (logical inconsistency) of A and $\neg A$. As is clearest in its explicit Tractarian form, the logicist aspiration is then to explain all genuine reason relations among sentences containing nonlogical terms by understanding some of them as

covertly having the form of logically complex sentences formed from other nonlogical expressions by applying logical connectives to them and deriving the good implications and incompatibilities that relate them as logically valid implications and logical inconsistencies.

Expressivism claims that if one thinks only about the relations between logical vocabulary and purely logical reasons, one will miss the most important relations between logic and reasons. For the expressivist, by contrast to the logicist, pure logic is a by-product of what is actually the characteristic use that marks some locutions as *logical* locutions. That characteristic, defining use is the application of those locutions to extend *non*logical base vocabularies, by adding logically complex sentences formed from the sentences of the base vocabulary by applying logical sentential connectives recursively to them. The relation between the base vocabulary and its logical extension is two-fold. The logically extended vocabulary is *elaborated from* the base vocabulary, and it is *explicative of* it. We abbreviate this two-fold functional specification of the expressive role characteristic of logical vocabulary by saying that it is "LX" for its base vocabulary: *el*aborated from and *ex*plicative of it.[2] The central idea of the "elaboration" side of this formula is that both the lexicon of sentences of the logically extended vocabulary and the reason relations of implication and incompatibility that articulate the conceptual roles of those sentences are uniquely and completely determined by and computable from the lexicon and reason relations of the base vocabulary that is logically extended. The central idea of the "explication" side of this formula is that the extension of the base vocabulary by the addition of logical vocabulary makes it possible to *say*, in the logically extended vocabulary, what implication and incompatibility relations hold not only in the base, but also in its logical extension.

These two sides of the functional characterization of the expressive role that defines logical vocabulary are not independent of one another. The explicative expressive function imposes an important criterion of adequacy on the rules for elaborating relations of implication and incompatibility for the logically extended vocabulary from the relations of implication and incompatibility that govern the material base vocabulary. The rules for introducing logical connectives—paradigmatically those that ensure that the conditional codifies implication and negation codifies incompatibility in accordance with the principles formulated above—must extend the relations of implication and incompatibility that they elaborate *conservatively*. That is, all the inferential relations among nonlogical vocabulary must be preserved, and no new implications or incompatibilities involving only nonlogical vocabulary must be introduced.

Why? Because introducing vocabulary to *express* those reason relations should not *change* them. It should *just* make it possible to express

explicitly in the (logically extended) object language the reason relations that *implicitly* govern the material base vocabulary (which we theorists express in a proof-theoretic *meta*language of sequents). We will see below that there are other notions of explication where it is not appropriate to impose a corresponding conservativeness condition. But conservativeness is required for the sort of explication that governs the introduction of specifically *logical* vocabulary.

A particularly vivid example of the trouble one can get into if one does not impose this condition is provided by the connective "tonk" introduced by Arthur Prior as part of an argument against natural deduction calculi without restrictions on the rules. He pointed out that the effect of using "tonk" with the introduction rule of classical disjunction, which includes:

$$\frac{p}{p \text{ tonk } q}$$

and the elimination rule of classical conjunction, which includes:

$$\frac{p \text{ tonk } q}{q}$$

is to license the implication:

$$\frac{\dfrac{p}{p \text{ tonk } q}}{q}$$

He called that a "runabout inference ticket," since it licenses the implication from arbitrary logically atomic *p* to arbitrary logically atomic *q*. Nuel Belnap diagnosed the trouble as a violation of conservativeness.[3] Introducing new logical vocabulary should not create any new implications involving only old vocabulary. Ignoring that constraint courts the risk of "tonking up" one's implication relation: trivializing it. The expressivist's requirement that logical vocabulary be explicative of the reason relations of the base vocabulary it extends offers a principled rationale for the requirement of conservativeness that Belnap imposed more or less *ad hoc*.

Although conservativeness is a criterion of adequacy (and so a necessary condition) for any vocabulary playing the distinctive explicative expressive role being recommended as demarcating specifically logical vocabulary, it is not appropriate to require it for nonlogical vocabulary in general. For it is characteristic of ordinary empirical descriptive vocabulary that its use involves endorsing nontrivial material implications relating its circumstances of appropriate application to its appropriate consequences of application. The nonlogical content of concepts such as copper, fever, and cruel incorporate many such substantive material implications

(and incompatibilities) that need not be redundant relative to other concepts antecedently available. This important dimension of content is ignored and made invisible by understanding conceptual content in terms of truth conditions. For that is the idea of conditions that are both individually necessary and jointly sufficient—the idea of circumstances and consequences of application that are guaranteed to coincide, so that no substantive inferential commitment is involved in the transition between them.

The covert ideology behind the notion of truth conditions is accordingly a kind of logicism. It is the idea that the reason relations that govern *all* vocabulary should not just be *expressible* by logical vocabulary (vocabulary governed by *logical* reason relations), but that those reason relations must be in the end *reducible to* logical reason relations. Since logical reason relations must be conservative over the rest of the language, so must reason relations in general. That is the origin of the requirement of the coincidence of circumstances and consequences of application in the form of conditions that are both necessary and sufficient. But that squeezes out nonlogical content, as articulated by *materially* good relations of implication and incompatibility: those that hold in virtue of the contents of *non*logical concepts. Here logicism about the relations between logic and reason relations misleads us about the latter, and does damage to our semantic theory (according to even a very weak form of semantic inferentialism).[4]

So the first half of the "LX" characterization of the expressive role characteristic of logical vocabulary is that the reason relations of implication and incompatibility that govern logically complex sentences must be conservatively elaborated from the reason relations governing the logically atomic sentences of some base vocabulary that is extended by the addition of those logically complex compounds of the atomic sentences. We use "vocabulary" here in a technical sense, as meaning a relational structure comprising a *lexicon*, which is a set of sentences, together with reason relations of implication and incompatibility defined on that lexicon. Syntactic formation rules must compute the lexicon of the logically extended vocabulary from that of the base vocabulary by the application of connective rules—which say, for instance, that if A and B are sentences in the extended lexicon, then so is $A \rightarrow B$. And an elaboration function must be specified that shows how to compute the reason relations governing sentences of the extended lexicon conservatively from reason relations governing sentences of the base vocabulary.

It is much less straightforward to say exactly what is meant by the portion of the LX-ness condition that requires the logically extended vocabulary that is elaborated from a base vocabulary to be *explicative of* the reason relations of the base vocabulary. We will have a lot more to say about this issue in the next chapter, when we introduce our

preferred universally LX logic, NMMS. The underlying thought is simple enough. For the expressivist, the paradigmatic logical connectives are the conditional and negation. The characteristic and defining expressive function of *conditionals* is to make explicit *implication* relations, and the characteristic and defining expressive function of *negation* is to make explicit *incompatibility* relations. Those connectives do that best if they satisfy these two conditions:

- Deduction-Detachment (DD) Condition on Conditionals:
  $\Gamma \mid\!\sim A \rightarrow B$ if and only if $\Gamma, A \mid\!\sim B$.
- Incoherence-Incompatibility (II) Condition on Negation:
  $\Gamma \mid\!\sim \neg A$ if and only if $\Gamma \# A$.

The first condition is inspired by Frank Ramsey's understanding of the expressive role of conditionals, and could with equal justice be called the "Dual Ramsey" Condition. It says that a premise-set implies a conditional just in case the result of adding the antecedent of that conditional to it is a premise-set that implies the consequent of the conditional. In this way, facts about implications are expressed in conditional sentences. Deduction-Detachment thereby offers one clear sense in which the conditional $A \rightarrow B$ *says that* A implies B. For if the underlying nonlogical base vocabulary contains the implication

It is raining $\mid\!\sim$ The streets will be wet

(which is materially a good implication in part because of the meanings of the non-logical terms "raining" and "wet"), and it does not contain the implication

The coin in my hand is copper $\mid\!\sim$ The coin in my hand would melt at 100°C.

then in the vocabulary that has been logically elaborated from that base, the conditional

It is raining $\rightarrow$ The streets will be wet

is provable or derivable from the first sequent, and there is no way so to become entitled to the conditional

The coin in my hand is copper $\rightarrow$ The coin in my hand would melt at 100°C.

in the absence of the second sequent. Neither one holds in virtue of logic alone, as does the more complex conditional

((It is raining → The streets will be wet) ∧ It is raining) → The streets will be wet.

which follows from Deduction-Detachment alone (in the context of Containment). The expressivist idea is that it is the fact that the nonlogical conditional can be justified (here, in a sequent calculus with sequents from the nonlogical, material base as leaves of the proof tree), and what it can in turn be used to justify that is the key to understanding the expressive role of conditionals. The fact that the more complex conditional holds just in virtue of the conceptual content of the conditional merely reflects that underlying expressive role.

The Incoherence-Incompatibility condition on negation says that a premise-set Γ implies a negated sentence just in case the sentence negated is incompatible with Γ. It offers one clear sense in which the negation $\neg A$ *says that* A is incompatible with a premise-set—namely, any premise-set that implies that negation. For it codifies an understanding of the negation of A as its minimal incompatible. $\neg A$ is implied by everything that is materially incompatible with A. The Aristotelian contradictory of A is what is implied by all the Aristotelian contraries of A. So <u>not-red</u> is implied by <u>green</u>, <u>blue</u>, and so on. In the logically extended vocabulary, one can say that "the swatch is green" is incompatible with "the swatch is red" by asserting the conditional with negated consequent

The swatch is green → ¬(The swatch is red).

Logicists want to explain material incompatibility (Aristotelian contrariety) in terms of logical negation—and so, inconsistency (Aristotelian contradiction). Expressivists seek to exploit these same conceptual connections in the other direction. (In the useful terms current in the philosophy of logic, the order of explanation that introduces the concept of logical negation on the basis of a prior concept of material incompatibility is understanding negation on the "Australian plan," by contrast to the contrary "American plan" (Berto and Restall, 2019).)[5]

Logical expressivists focus on the directly *expressive* sentential logical connectives: conditionals and negation. These contrast with merely *aggregative* logical connectives, paradigmatically conjunction and disjunction. The function of these auxiliary connectives is to make explicit what is indicated by commas on the two sides of the turnstile in a sequent-calculus formulation of reason relations. We sometimes refer derisively to the merely aggregative connectives as "Boolean helper-monkeys."

Chapter Three shows how the expressivist aspirations sketched here in terms of the LX-ness of logic for arbitrary base vocabularies can be made more definite, and then how those aspirations can be realized in the familiar form of a multisuccedent sequent calculus. We also show there how the logic we endorse can be enriched with modal operators whose expressive role is to make explicit features of the local and global *structure* of the reason relations that are in turn made explicit by ordinary, nonmodal sentential connectives.

Thinking about the relations between reasons and logic according to the expressivist model being recommended here offers a distinctive perspective on traditional philosophical notions of rational consciousness and self-consciousness thematized most prominently by classical German idealists. Adopting practical discursive attitudes of accepting and rejecting claimables, which stand to one another in reason relations of implication and incompatibility, defines a distinctive sense in which creatures can qualify as *conscious*. This is consciousness in the sense of sapience or apperception, rather than sentience or perception. That sort of conceptually articulated awareness is not just a matter of being awake and sensorily responsive—able, say, to feel pain and distinguish red things from green ones. It consists rather in taking it *that* things are thus-and-so, in having propositionally contentful attitudes expressible by what count as declarative sentences just in virtue of expressing contents that can stand to one another in material reason relations of implication and incompatibility—to serve as and stand in need of reasons. Apperceptive consciousness in this sense is accordingly to be understood as *rational* consciousness: consciousness articulated by reasons.

Understanding the relation of logic to those reason relations as the logical expressivist does then presents logical vocabulary as the organ of a distinctive kind of rational *self*-consciousness. For it enables the making explicit in claimable form of the reason relations that articulate rational, apperceptive consciousness. Chapter Three shows how this sense of specifically *logical* rational self-consciousness as logical explicitation can be made precise in the form of a powerful expressive completeness theorem (tied to what we call "explicitation by sequents") relating what can be said in the logical extension of a base vocabulary to the reason relations of that base vocabulary.

The two complementary approaches to the reasons question about logic, logicism about reasons and rational expressivism about logic, agree in one feature of their ambition: its generality. Logicism seeks to exhibit *all* (doxastic, nonpractical) reasons as covertly *logical* reasons—good, if and insofar as they are good reasons, in virtue of being reducible to logically valid implications (in the case of reasons-for) and logical inconsistencies (in the case of reasons-against). Analogously, the rational expressivists'

ideal logic would be LX for *every* base vocabulary. The connective rules defining its logical vocabulary would permit it conservatively to extend, elaborate its reason relations from, and articulate the conceptual contents of, logically complex sentences that express the reason relations of *any* base vocabulary whatsoever. Expressivists can recognize vocabularies with more limited expressive power as genuinely logical—for instance, logics that can codify only mathematical implications, or (more to the point for our larger argument here) only structurally closed, monotonic and transitive reason relations. Insofar as they are LX for some restricted range of base vocabularies, they can still qualify as proper logics for the logical expressivist. This is the sense in which there can be logics of colors, or sortals, or quantum phenomena—just insofar as the vocabulary in question can be elaborated from and is explicative of color vocabularies, sortal terms, or the vocabulary of quantum mechanics. But the expressivist ideal is expressive *universality*: a logic that is LX for any vocabulary whatsoever, no matter what its lexicon or reason relations might be. The ultimate measure of a logic—the normative standard for assessments of its success *as* a logic—is its expressive power, the range of reason relations it can make explicit. Perhaps universal LX-ness is unobtainable, practically, or in principle. But anything less counts as expressive impoverishment.

From the point of view of rational expressivism about logic, this ideal shows up as what is right about traditions in the philosophy of logic that pick out purely logical reason relations as those that are topic neutral and (so) maximally general. Logical reasoning, the thought goes, applies to any and every subject matter, by contrast to content-specific reason relations that hold only for some topics. One way of articulating this thought originates with Bolzano, is exploited by Frege, and is given voice by Quine. It is a methodology of noting invariance of reason relations under substitution.[6] An implication or incompatibility holds in virtue of the logical form of the sentences involved just in case it is robust under the (uniform) substitution of nonlogical for nonlogical vocabulary. Considering reason relations that are invariant under arbitrary substitution of nonlogical for nonlogical vocabulary is a clear and precise way of operationalizing the idea that the reason relations being assessed hold independently of the nonlogical subject-matter of the sentences involved.

This is an idea logical expressivists can wholly endorse. For notice that it presupposes a notion of the goodness of *material* (nonlogical) implications and incompatibilities. It picks out some reason relations as good in virtue of their purely logical form in case two conditions hold:

i) They are good, and

ii) All of their substitutional variants, resulting from replacement of nonlogical by nonlogical vocabulary, are also good.

(ii) is the condition that the implications and incompatibilities be good under arbitrary substitutions *salva consequentia*. Both conditions appeal to a notion of the goodness of an implication or incompatibility relation that is *not* a matter of their *logical* goodness (that is what is being defined), and must be intelligible in advance of the notion of formal logically valid implications and inconsistencies, so as to be available for use in the definition. Of course, for the expressivist about logic, the procedure only works if one has available a vocabulary that includes logically complex sentences formed from the sentences of some nonlogical base vocabulary—that is, vocabularies that result from the logical elaboration of those base vocabularies. The dependence of this way of understanding and justifying the universality of logic on prior material relations of implication and incompatibility puts it in substantial tension with the logicist's aspiration to understand those relations in terms of implications and inconsistencies that hold in virtue of logic alone.

It is worth noticing in passing that the Bolzano-Frege-Quine method of noting invariance under substitution is really a way of thinking about the concept of form generally, rather than a way of thinking about *logic* and *logical* form specifically. For it does not at all depend on the vocabulary that is kept invariant being logical vocabulary. Any vocabulary at all can be substitutionally privileged for these purposes. If one looks at materially good implications that remain good under arbitrary substitution of non-*theo*logical for non-*theo*logical (or non-geological for non-geological, or non-culinary for non-culinary, non-nautical for non-nautical …) vocabulary, the result will be implications that hold in virtue of their theological (geological, culinary, nautical …) form alone. (These might relate terms such as "sacred" and "profane", "sin" and "peccable," and so on.) In particular, when the substitutional methodology is applied to logic as supplying the invariance-privileged vocabulary, it presupposes a distinction between logical and non-logical vocabulary. That methodology accordingly cannot be appealed to in demarcating logical vocabulary. Our expressivist proposal for that demarcational task is that logical vocabulary is vocabulary that is LX for some base vocabulary—in the ideal limit, *universally* LX, in being capable of being elaborated from and explicative of any and all base vocabularies.

2.4 The Structure of Reason Relations

The logical expressivist aspiration to universal explicative generality, in the sense of seeking a logic that can codify any and all reason relations, has one consequence that is of particular significance for the logical and semantic innovations that we will propose here. That is that an ideal logic would not build in any undue restrictions on the *structure* of the relations

of implication and incompatibility that it can express. This observation points to a fundamental issue that is in general not accorded the centrality in contemporary philosophy of logic that we believe it deserves. There is a mismatch between the structure of specifically *logical* reason relations and the structure of reason relations in general.

For most of the last century it has been agreed (following Tarski (1936) and Gentzen (1934)) that logical consequence has the topological structure of a closure operator. That means that for any premise-set X (a set of sentences drawn from some lexicon), it makes sense to talk about *the* set $c(X)$ of its logical consequences. That consequence-set *c(X)* is the *closure* of the premise-set X under logical consequence in the sense that it meets three structural conditions:

- The premises are included in the consequence set:
  $X \subseteq c(X)$. (Containment)
- Adding further premises does not lose any consequences:
  $c(X) \subseteq c(X \cup Y)$. (Monotonicity)
- The consequences of the consequences of any premise-set are already consequences of the original premise-set:
  $c(c(X)) = c(X)$. (Idempotence, entailing Transitivity)

The consequence relations of classical and intuitionist logics, and the modal logics built on them such as S4 and S5, as well as traditional multivalued logics such as Strong Kleene (K3) and its dual, Priest's Logic of Paradox (LP) are all structurally closed in this topological sense (defined in topology by the Kuratowski closure axioms, which inspired Tarski's metalogical application of them).

The implication relations that govern most of our practices of challenging and defending claims do not have this closure structure, though. From the point of view of traditional logic, ordinary material non- or prelogical reasoning (as well as its sophisticated institutionalized versions in law, medicine, and the special sciences) is in general *sub*structural. We will say that it has an *open* structure, by contrast to the *closed* structure of traditional logics. The substructurality of reasoning in general is most obvious in the failure of monotonicity. Almost always when we give reasons for or against our own assertions or those of others, those reasons are defeasible. That is, there are further considerations that, if admitted as auxiliary hypotheses or collateral premises, would infirm the implication or incompatibility connecting the proferred reasons to what they are presented as reasons for or against. We take such defeasibility for granted in casual justifications of or challenges to quotidian claims. If someone argues "Everyone is already in their offices, so we will be able to start the

meeting on time," they know full well that a myriad of interfering events—last-minute telephone calls, bouts of forgetfulness, fire alarms …—could infirm the implication. But more serious scholarly and scientific arguments are like this, too. The carbon 14 dating results provide good reasons for dating the White Sands footprints to 20,000 years ago—but not if the association of the footprints with the *Ruppia* seeds that were actually dated is spurious, or if the hard-water effect had incorporated more ancient carbon atoms into those seeds, and so on.[7] The reasons offered supporting or contesting medical diagnoses or clinical predictions are understood to be overrideable by all sorts of possible further evidence. The same is true of legal argumentation, where distinctive forms of defeasibility of reasoning are institutionalized in the form of appeals courts.

As was pointed out in the Introduction, two broad and important areas of philosophical thought about reasoning where the nonmonotonicity of implications is explicitly acknowledged, thematized, and theorized about are probabilistic reasoning and implications codified in subjunctive conditionals. An argument that justifies with high probability a certain conclusion about a population given one specification of the reference class from which the sample is drawn can turn into a much weaker argument, or even an argument with equally high probability for a contrary conclusion, if further information is added that narrows the reference class. Adding further premises can flip the valence of subjunctive reasoning:

> If I were to press this key, the program would start.
> If I were to press this key and Dan rewrote the code last night to require a different key, then the program would not start.
> If I were to press this key and Dan rewrote the code last night to require a different key, but Ulf noticed this morning and switched it back, then the program would start.
> But not if, in addition, the computer is unplugged.
> Unless it has a battery backup.
> But even then, if the battery has run down, the program won't start.
> …

One might be tempted to treat all of these familiar phenomena as reflecting low pragmatic considerations of conversational convenience, rather than deep structural features of implications and incompatibilities. The thought would be that it is simply too laborious and tedious to put in all the qualifications required for fully stating a *real* reason—that is, an indefeasible, dispositive reason. So one just gestures in the direction of the complicated, extensive conjunction of evidentially significant premises or conditions that is implicitly being invoked when reasons are offered. But that is not right. For in the sorts of cases being considered, there typically

is no totality of potentially relevant considerations, no full specifications of the further circumstances under which an implication could fail. Any finite list will be found to have left out some outré possibilities. The hungry lioness notices the nearby wounded gazelle, so she will chase it. But not if she is paralyzed, struck by lightning, suddenly encased in glass, reduced to microscopic size, teleported to Venus, and so on. And there is in general no nontrivial way to specify an infinite list that would be comprehensive. By "nontrivial" here is meant any way that is not equivalent to saying "The hungry lioness notices the nearby wounded gazelle, so she will chase it—except if for some reason she does not." If "all other things being equal" qualifications had the effect of ruling out all potentially defeating further premises, then such *ceteris paribus* clauses would be trivial in this objectionable way. In fact they should not be understood as somehow *removing* defeasibility and rendering implications monotonic. (A Latin phrase whose utterance can do *that* is properly called a "magic spell.") The real expressive function of *ceteris paribus* clauses is just to acknowledge explicitly that the implication being endorsed *is* defeasible. Such a concession is not a taking-back of the argument to which it is appended. Rather, it pragmatically marks the openness of a conversational pathway in which some of the open-ended set of possible defeaters are actually asserted, and so themselves become subject to critical rational challenge and defense.

Note that monotonicity is as well-defined for incompatibility as for implication. An incompatibility property is nonmonotonic if adding elements can make incompatible sets compatible. And monotonicity is equally implausible as a constraint on material incompatibility relations. Indeed, there is a general procedure for turning failures of monotonicity for material consequence relations into failure of monotonicity for material incompatibility relations. The nonmonotonicity of the implications codified in subjunctive conditionals is illustrated most forcefully by the possibility of "Sobel sequences," in which further information sequentially flips the valence of the implication. If I were to strike this dry, well-made match, it would light. But *not* if it is in a very strong magnetic field. Unless, additionally, it were in a Faraday cage, in which case it would light. But *not* if the room were evacuated of oxygen. And so on. But here we can see that reasons *against* a claim are as defeasible in principle as reasons *for* a claim. Striking the dry, well-made match and its lighting is incompatible with its being in a strong magnetic field. But striking the dry, well-made match and its lighting and its being in a Faraday cage *is* compatible with its being in a strong magnetic field. Negation operators introduced to codify material incompatibility relations will have to deal with the structural nonmonotonicity of incompatibility every bit as much as conditionals introduced to codify material consequence relations will

have to deal with the structural nonmonotonicity of implication. That this observation about the open structure of both kinds of reason relation is of some importance becomes clear in light of the fact that prominent approaches to nonmonotonic logics (such as preferential models and default reasoning) help themselves to classical negation and specifically, the monotonic relation of inconsistency it supports (Kraus et al., 1990; Horty, 2007).[8]

To say that material relations of implication and incompatibility are not in general monotonic is to say that they are not *always* monotonic, not that they are *never* monotonic. The nonlogical implication of "Monochromatic swatch #13 is red," by "Monochromatic swatch #13 is crimson," and the nonlogical incompatibility of both with "Monochromatic swatch #13 is green," continue to hold when arbitrary further collateral premises are added. The same is true for many implications and incompatibilities involving claims about Euclidean plane figures, organisms classified according to Linnaean nomenclature, minerals specified by their chemical composition, and a host of other subject-matter specific reason relations.[9] This is what made it tempting and plausible for mid-twentieth century philosophers to talk about the "logic" of color-talk and other vocabularies that are not themselves logical vocabularies. Understanding ordinary empirical descriptive concepts (among others) involves not only having some practical ability (however partial and fallible) to distinguish good from bad implications and incompatibilities governing their possible applications, but also having some practical ability (however partial and fallible) to associate with them ranges of subjunctive robustness. That is the capacity to distinguish additional circumstances (premises) that would not defeat the implication or ruling out of the conclusion, from those that would. In navigating the complexly structured constellation of reasons for and reasons against in the course of defending and challenging various nonlogical claimables, the cases where the ranges of subjunctive robustness of the reason relations are *total* provide essential landmarks by which discursive practitioners can orient themselves. These are local regions of *persistent* implications and incompatibilities: ones that can be counted upon to remain good in the face of further collateral information or commitment.

The fact that these islands of monotonicity (or transitivity) of material reason relations are in general surrounded by seas of nonmonotonic implications and incompatibilities suggests another dimension of the expressive task of making explicit reason relations that expressivists about logic might reasonably call on logical vocabulary to perform. In addition to being able to use conditionals and negated sentences of a logically extended vocabulary to codify particular implications and incompatibilities that hold in the material base vocabulary to which logical vocabulary has been added

(and from which its reason relations have been elaborated), it would be good to be able explicitly to mark regions of such reason relations where various elements of closure structure hold locally, even though they do not hold globally. In the next chapter, we will show how this can be done quite generally. For now it is enough to introduce the general idea.

Suppose we mark persistence of a particular implication by an upward arrow attached to the snake turnstile. If it is not only the case that $\Gamma \mid\!\sim A$, but also that for any and all sets of sentences Δ, $\Gamma \cup \Delta \mid\!\sim A$, then we can express that by writing $\Gamma \mid\!\sim^{\uparrow} A$. That local fact about the structure of implication in the base vocabulary can then be made explicit in a logically extended vocabulary if we add a new bit of logical vocabulary with a definition parallel to those used to illustrate the desired expressive role of the conditional and negation:

- Persistence-Codifying Modal Operator:
  $\Gamma \mid\!\sim \Box A$ if and only if $\Gamma \mid\!\sim^{\uparrow} A$.

If Γ does not just imply A, but persistently implies A, then we say that it also implies $\Box A$. That symbol is chosen because it can be read as "necessarily"—in the sense of "come what may," or "no matter what else is true." It explicitly marks the fact that the range of subjunctive robustness of the implication is *total*. The implication continues to hold no matter what additional premises are added to it. The box explicitly marks the monotonicity of the implication.

It is worth looking a little more closely at the sense in which the job of expressing explicitly the local structural feature of persistence or monotonicity of particular implications can be done by introducing logical vocabulary with a distinctively modal flavor. For it brings into view an important sense in which the corresponding principle DD we have suggested for thinking about the expressive role of conditionals defines them, too, as *intensional* connectives. If we fix the lexicon of a base vocabulary, we can think of the set of sets of sentences of that lexicon as a universe of points of semantic evaluation, playing a role analogous to that of possible worlds. Each such potential premise-set can be thought of as surrounded by a cloud of further sentences: those that are implicitly included in or excluded by that premise-set, in the sense of being implied by or incompatible with it: $\{X : \Gamma \mid\!\sim X\}$ and $\{Y : \Gamma \# Y\}$. The logical connective $\rightarrow$ is intensional in that the rule for introducing conditional sentences says that in order to determine whether or not $\Gamma \mid\!\sim A \rightarrow B$, one cannot just look at what is "true" or "false" at the point of evaluation Γ—in the sense of what is implied by or incompatible with Γ. Rather, one must look at what is "true" or "false," implied by or incompatible with, another, different, neighboring point of semantic evaluation, namely $\Gamma \cup \{A\}$. If B is "true at"

or implied by that premise-set, then the logically complex sentence $A \to B$ is implied by Γ. In one core usage, the difference between extensional and intensional vocabulary is whether the status (paradigmatically, truth) at one point of semantic evaluation of sentences formed using that vocabulary is a function just of that same point of semantic evaluation, or whether to compute that semantic value at one point of evaluation, one must consult other points of semantic evaluation. The monotonicity-codifying box operator is (like a Ramsey-test conditional satisfying the DD condition) an intensional operator in this sense. For whether or not $\Gamma \mid\!\sim \Box A$ depends on whether A is a consequence, not just of Γ, but also of all the supersets of Γ. This modal operator illustrates the principle that logical vocabulary can be called on to make explicit not just which particular reason relations a given premise-set stands in, but also facts about the structural conditions that hold locally, in its neighborhood, even if they do not hold globally.

Material incompatibility, we have argued, is symmetric—unlike implication. $\Gamma, A\#B$ if and only if $\Gamma, B\#A$. It follows that facts about the incompatibility of sentences with premise-sets are equivalent to facts about the incoherence of premise-sets. The two incompatibility facts just instanced are equivalent not only to each other, but to the incoherence of the set $\Gamma \cup \{A\} \cup \{B\}$. We can introduce a special symbol, $\perp$ (pronounced "perp") appearing only (alone) on the right of the turnstile, marking the set that appears on the right as incoherent: $\Gamma, A, B \mid\!\sim \perp$. The same effect is achieved by using an empty right-hand side: $\Gamma, A, B \mid\!\sim$. In either case, incompatibilities are traded for the incoherence of sets, and the incoherence of sets is expressed using the turnstile of implication. This common notational convenience has tended to lead—as a matter of psychology rather than philosophy—to a failure to acknowledge the co-equal status of incompatibility with implication as a basic form of reason relation. That is why we have, until now, instead used the nonstandard "#" notation to express incompatibility. In the fond hope that the temptation to misunderstand the significance of the notational convenience of expressing material incoherence of sets with the snake turnstile of nonmonotonic implication can be resisted, we will henceforth sometimes avail ourselves of that facility.

The fact that material incompatibility is not always monotonic, then, is another way of saying that material incoherence of sets of claimables is not monotonic. The fact that Γ is incoherent does not entail that $\Gamma \cup \Delta$ is incoherent. In the general case, rational incoherence (incompatibility) need not be persistent. If in some case it is not, if $\Gamma \mid\!\sim \perp$ and it is not the case that $\Gamma, \Delta \mid\!\sim \perp$, then the incoherence is curable, or defeasible. "Tweety is a bird" and "Tweety cannot fly" are incompatible, but adding "Tweety is a penguin" restores coherence. For a more serious example, wave behavior and particle behavior are incompatible in classical mechanics. In quantum mechanics, with lots of other auxiliary hypotheses added,

they become compatible. Of course, *some* incoherences *are* persistent, in that not only is Γ incoherent, but so are all of its supersets. The combination of the claim that plane figure *T* is a Euclidean triangle and the claim that the angles of *T* add up to more than two right angles is incoherent (the claims are incompatible). And no further premises can cure that incoherence. Traditional logics, which treat both implication and incompatibility/incoherence as having full topological closure structure (so, as monotonic), enshrine a particularly striking relationship between the two reason relations. This is the principle "*ex contradictione quodlibet*" or explosion: from a contradiction, everything follows. As a structural principle regarding material reason relations of implication and incompatibility/incoherence generally, it is hard to know what to say in favor of this principle. It does not seem to correspond to anything in our ordinary reasoning practices, nor their institutionalized forms in scholarship, engineering, theoretical science, clinical medicine, or the law. This is a point on which structural logicism clearly seems to lead us astray. (That is why this principle is a permanent source of classroom embarrassment to teachers of elementary logic.)

However, the distinction between *persistently* incoherent premise-sets and those that are only curably or defeasibly incoherent is invisible to traditional logic, because of its commitment to the monotonicity of reason relations. Awareness of the distinction opens up the possibility of acknowledging that for ordinary, nonpersistently incoherent sets, we can still distinguish—as indeed, we do—between what follows from them and what does not. One might maintain that distinction even in cases where one does not have any idea what further premises might in principle restore coherence, were they added to the premises actually on hand. The fact that traditional logic offers no help in understanding this fact about them should not blind us to the possibility and indeed actuality of material base vocabularies in which at least potentially curably incoherent premise-sets do not explode implicationally. Expressivists about logic should accordingly aspire to construct logics in which incoherent premise-sets (except perhaps the incorrigibly, that is, persistently, incoherent—for instance, logically inconsistent—ones) are as tractable as coherent ones are. In Chapter Three we will see how this can be done—and further, how local regions of incoherence of premise-sets can explicitly be marked in logically extended base vocabularies where not all incoherence is incorrigible.

We can say a bit more about the relations among the different components of topological closure structure. It is widely recognized that failures of monotonicity (MO) generate failures of simple transitivity. So, a standard example of an MO failure is:

- Tweety is a bird, so Tweety can fly.

But *not*

- *Tweety is a bird and Tweety is a penguin, so Tweety can fly.

The corresponding failure of simple transitivity is

- Tweety is a penguin, so Tweety is a bird.
- Tweety is a bird, so Tweety can fly.

But *not*

- *Tweety is a penguin, so Tweety can fly.

Ryan Simonelli of our ROLE group has observed further that there is a general procedure for turning examples of failures of monotonicity into examples of failures of *cumulative* transitivity (CT): cases where explicitating a consequence adds further consequences. It is *not* the case that

- *Tweety is a bird, so Tweety is a non-penguin.

But we do have

- Tweety is a bird, so Tweety can fly.
- Tweety is a bird and Tweety can fly, so Tweety is a non-penguin.

Adding the consequence (fly) to the premise-set (bird) that does not imply non-penguin yields a premise-set that *does* imply non-penguin. In this way we can see that where one finds nonmonotonicity in material consequence relations, one will also find failures of CT.

There is a further argument against imposing transitivity as a global structural constraint. This is the observation that CT (in the context of CO), together with an implication-codifying (double) Ramsey conditional—that is, one that satisfies the Deduction-Detachment condition—forces MO (Hlobil, 2016). For if we start with some arbitrary implication $\Gamma \mathrel{|\!\sim} A$, we can derive $\Gamma, B \mathrel{|\!\sim} A$ for arbitrary B—that is, we can show that arbitrary additions to the premise-set, arbitrary weakenings of the implication, preserves those implications. And that is just monotonicity. For we can start with $\Gamma \mathrel{|\!\sim} A$ and the CO-instance $\Gamma, A, B \mathrel{|\!\sim} A$ and argue:

$$\dfrac{\dfrac{\Gamma \mathrel{|\!\sim} A \qquad \dfrac{\Gamma, A, B \mathrel{|\!\sim} A}{\Gamma, A \mathrel{|\!\sim} B \to A}\,[\text{Ramsey Condition Right-to-Left}]}{\Gamma \mathrel{|\!\sim} B \to A}\,[\text{CT}]}{\Gamma, B \mathrel{|\!\sim} A}\,[\text{Ramsey Condition Left-to-Right}]$$

As a result, a proper conditional (one that satisfies the Deduction-Detachment condition) cannot be introduced conservatively on a nonmonotonic base. Such a conditional cannot in principle *explicate* (conservatively express) a nonmonotonic base. If we want such a conditional (and CO), we must forego CT as a global principle. For it is not just that CT creates a problem (forces monotonicity) if the language already contains an implication-codifying conditional. The problem is that we cannot *add* such a conditional to a nonmonotonic base language without endorsing new (monotonic) implications involving only the old, prelogical vocabulary. And that violates the explicative, implication-codifying expressive task characteristic of conditionals.

The idea that reason relations in general do, or even must, share the topological closure structure of specifically *logical* reason relations is a restricted kind of logicism: *structural* logicism about reasons. The nonmonotonicity (and, as we will see, also nontransitivity) of nonlogical implication relations—the fact that they are not necessarily or in general monotonic or transitive—shows that the structural logicist thesis is false. That fact poses a challenge to the more general logicist program of reconstructing all good reasoning on the basis of logically good reasoning: logically valid implications and logical inconsistencies. If structural logicism about material reason relations is not true, then neither is logicism about reasons generally—at least not in its original form. That challenge need not be understood as fatal to the general logicist program. Perhaps nonmonotonic implications can somehow be reconstructed out of monotonic ones. Something like that is what traditional nonmonotonic logics aspire to: to start with proper, monotonic logical systems, and tweak them somehow (add structure such as preference orderings, distinguish different classes of premises, restrict models) so as to yield nonmonotonic reason relations.

But this structural difference between logical reason relations and reason relations more generally presents a problem for expressivism about logic just as much as it does for logicism about reasons. For it is not at all obvious how monotonic logical vocabulary could make explicit nonmonotonic implications and defeasible incompatibilities. Expressivists have a project here that is different from that of traditional nonmonotonic logicians. The important task is not seen to be constructing nonmonotonic *logics*, but rather logically codifying nonmonotonic (and nontransitive) consequence relations (and incompatibilities). It turns out, as we'll see in the next chapter, that to do that one does *not* need logics that are themselves nonmonotonic (or nontransitive). Though it sounds paradoxical, logics that are expressively adequate to make explicit radically substructural, open material reason relations can themselves exhibit the full topological closure structure characteristic of classical logic.

To see how this can be possible, one must distinguish two different notions of consequence (and incompatibility) that govern the use of logical vocabulary. First are the implication relations that hold among sentences of some particular logically extended base vocabulary. If the reason relations of the base vocabulary are themselves structurally open (nonmonotonic or nontransitive), as we have argued they generally are, then since the reason relations of the logically extended vocabulary elaborated from them must be conservative, those reason relations, too, must be structurally open. If $p \mid\!\sim q$ holds in the base vocabulary and $p, r \mid\!\sim q$ does not hold in the base vocabulary, then $\mid\!\sim p \rightarrow q$ but not $\mid\!\sim (p \wedge r) \rightarrow q$ will hold in the logical extension of *that* base vocabulary. Note that in nonmonotonic settings, the fact that $p \rightarrow q$ is implied by the empty set of premises, and is in that traditional sense a theorem, does *not* entail that it is implied by *every* set of premises. For instance in the situation being envisaged, the fact that $\{p, r\}$ does not imply q does not entail that r does not imply $p \rightarrow q$. The notion of "theorem" that matters in nonmonotonic settings is being implied by *every* premise-set, rather than being implied by the empty premise-set. Global structural monotonicity runs these two distinct notions together.

But that failure of monotonicity in the logical extension of some base vocabularies is entirely compatible with the *purely logical* implications, the implications that hold in virtue of logical form alone, being fully monotonic. For the sequent codifying a nonmonotonic implication is derivable only in proof trees that have as leaves sequents of the particular base vocabulary in which p materially implies q and $\{p, r\}$ does not materially imply q. Purely logical implications, implications that hold in virtue of logic alone—that is, the connective rules (plus Containment)—hold regardless of what base vocabulary one uses those connective definitions to extend, including those in which p does not materially imply q and $\{p, r\}$ does materially imply q. The expressive logic we will recommend, NMMS (and, indeed, the even more expressively capable non-contractive version of NMMS), is capable of making explicit arbitrary, radically substructural base vocabularies: those that are in general neither monotonic nor transitive. But it is itself not a nonmonotonic or nontransitive logic. It is not only supraclassical, there is a sense in which it just *is* (a formulation of) classical logic.

2.5 Even More Radically Open Structures of Reason Relations

So relations of consequence (following from or being a reason for) in general do not have the full structure of closure operations—as the tradition of Tarski and Gentzen takes it that specifically *logical* consequence relations not only do, but must. In the same way, and for the same reasons, material relations of incompatibility are not in general monotonic. In view of these

features of ordinary, nonlogical reasoning, expressivists about logic should aspire to codify open-structured reason relations, including nonmonotonic and nontransitive ones. That is to give up *structural* logicism, and to remove much of the ground supporting logicism about reasons generally. One question that then arises is whether there are any structural conditions weaker than the full topological closure structure that *can* be counted on to characterize material reason relations.

Containment (CO) says that in the case of consequences of implications that are also premises of those implications, one can weaken the premise-set with arbitrary additional premises without infirming the conclusion. It is plausible that among the consequences of any premise-set are those premises themselves. Relevance logicians, however, object. They do not recognize as good, implications that involve premises that are irrelevant to the conclusions. That is a reason to treat the addition of irrelevant premises as able to defeat otherwise even trivially good implications such as instances of Reflexivity, of the form $A \mid\!\sim A$. Whether or not the relevantists are right about that, logical expressivists ought to aim for a logic that does not build in the presupposition that they are wrong. Put another way, a logic that could codify even consequence relations that fail to satisfy Containment is clearly more expressively powerful (other things being equal) than, and so preferable to, one that cannot. The logic NMMS introduced in Chapter Three works just fine (is provably LX) even for base vocabularies whose consequence relations do not satisfy Containment—though, as we will see, the result is that the logical consequence relations defined are no longer supraclassical.

Denying Monotonicity (MO) means not assuming that one can weaken *every* implication with arbitrary additional premises without infirming the conclusion. Is there a way to specify some more restricted set of additional premises one could add that *would* be guaranteed not to defeat an implication? One plausible candidate answer to this question is given by what has been called "*cautious* monotonicity." Cautious monotonicity says that while one might indeed not be able to weaken an implication by adding just any sentence as a collateral premise without running the risk of infirming it, it should at least be safe to add further premises that are already implied by the original premise-set.

Cautious Monotonicity (CM):

$$\frac{\Gamma \mid\!\sim A \qquad \Gamma \mid\!\sim B}{\Gamma, A \mid\!\sim B}$$

The idea is that since the premise-set Γ already implies A, adding A as an explicit premise should not cause any trouble with other consequences of Γ. Even though there might be *some* additional premises that *would* infirm

the implication, sentences that are *already implied* by the premise-set are not among them.

It has often been argued not only that cautious monotonicity is a plausible principle, but that it is in effect indispensable: that it is a *minimal* condition that well-behaved nonmonotonic consequence relations must satisfy.[10] Satisfying CM is generally regarded as a criterion of adequacy for assessing nonmonotonic logics. CM plays a prominent role, for instance in what Kraus, Lehman, and Magidor call the "core properties" or the "conservative core" of non-monotonic systems (and for this reason are now often called the "KLM structural properties" required of nonomonotonic systems), and count it a signal virtue of their preferential semantics for nonmonotonic logic that it validates this structural principle (Kraus et al., 1990). By contrast to monotonicity, there has not been much skeptical philosophical attention directed at cautious monotonicity. This is a shame, because the underlying issues here are just as important, and addressing them is deeply revealing of considerations that remain invisible if the discussion remains at the level of the much stronger structural principle of monotonicity.

The first step in appreciating this is realizing that cautious monotonicity is the dual of cumulative transitivity, a version (the shared context version) of Gentzen's "Cut."[11] This structural principle is expressed in Tarski's algebraic metalanguage for consequence relations by the requirement that the consequences of the consequences of a premise-set are just the consequences of that premise-set, and by Gentzen as the principle that adding to the explicit premises of a premise-set something that is already part of its implicit content does not add to what is implied by that premise-set. It is the principle appealed to in chaining together implications in extended consecutive reasoning.

Cumulative Transitivity (CT):

$$\frac{\Gamma \mathrel{|\!\sim} A \qquad \Gamma, A \mathrel{|\!\sim} B}{\Gamma \mathrel{|\!\sim} B}$$

CT says that adding a consequence of a premise-set to that premise-set never *adds* consequences—that what a premise-set implies when we add to it its own consequences, it already implies all on its own—while CM says that adding consequences to the premise-set never *subtracts* consequences the original premise-set had.

Here is a way to think about the underlying issue. Using language that was second nature to Leibniz and Kant, we can think about the *content* of a set of claimables in the literal sense of what is *contained in* it. A premise-set $\Gamma = \{A_1, A_2, \ldots, A_n\}$ literally contains all of the sentences A_i in the set-theoretic sense that these are the elements of the set Γ. We may say that

it contains them *explicitly*, since they are what we specify when we specify the set. As so contained, they are the *explicit content* of the set—"content" etymologically being what is contained. If it now happens that Γ implies *A*—in our notation, $\Gamma \mid\!\sim A$—then we can say that *A* is *implicit* in Γ, in the literal sense of being *implied by* it. *A*, then, is part of the *implicit content* of the premise-set Γ. (Analogously, we might think of every set Δ that is materially *incompatible* with Γ as being part of Γ's implicit *contrastive* content.)

Then CM and CT can be thought of as having a common topic. Both concern what happens when the status of some consequence of a premise-set is changed, by turning it into an additional premise. The process of moving a sentence from the right-hand side of the implication turnstile to the left-hand side, from appearing as a conclusion to appearing as a premise, might be called the process of *explicitation*. For it is the process of making some *implicit* bit of content *explicit*, turning what is *implicitly* contained in a premise-set (that is, what is implied by it) into a premise that is explicitly contained in the (new) premise-set.[12] Explicitation in this sense is not at all a *psychological* matter. And it is not even yet a strictly *logical* notion. For even *before* logical vocabulary has been introduced, we can make sense of explicitation in terms of the structure of non-logical, *material* consequence relations. Noting the effects on implicit content of adding as an explicit premise sentences that were already implied is already a process available for investigation at the methodological level of the *prelogic*, in the base vocabulary, before that material base is logically extended.

Both CT and CM concern the effects that explicitation has on the *consequences* of the premise-set, comparing the consequences before explicitation with the consequences after explicitation. Since the consequences of a premise-set are its implicit content, CT says that explicitation does not *gain* any implicit content, and CM says that explicitation does not *lose* any implicit content. CT says no consequences are added, and CM says no consequences are subtracted. Together, they entail that ***explicitation is inconsequential***: making explicit some or all of the implicit content of a premise-set has no effect on its consequences at all. Moving a sentence from the right-hand side of the implication-turnstile to the left-hand side does not change the consequences of the premise-set. It makes no difference whatever to the implicit content, to what is implied.

We began by opposing an *expressivist* approach to understanding the relations between logic and reasoning (mediated by reason relations of implication and incompatibility) to a *logicist* approach to understanding them, according to which all good reasons are *logically* good reasons: every genuine implication is valid in virtue of the logical form of its premises and conclusion. We then considered a weaker, purely *structural* form of logicism. It claims that the algebraic structure of material reason relations

of implication and incompatibility is the same as the algebraic structure of specifically *logical* relations of implication and incompatibility. For historical reasons we have gestured at, philosophers of logic have taken that algebraic structure to be topological closure. Topological closure is a matter of satisfying monotonicity and transitivity, MO and CT (as well as Containment, CO, or at least Reflexivity, RE). The present claim is that it is more philosophically revealing to focus on a different kind of closure structure, which involves pairing CT not with MO, but with CM. This might be called "*explicitation closure*," since it entails that explicitation is inconsequential. Since MO entails CM, rejecting the explicitation-closure form of structural logicism—by denying that explicitation is inconsequential for material consequence relations—will entail rejecting the topological-closure form of structural logicism. Consequence relations for which CT fails are *nontransitive*. Implication and incompatibility relations for which MO fails are *nonmonotonic*. Reason relations for which CM fails—implications and incompatibilities can be defeated even by the addition of premises that are already implied by a premise-set—are even more radically substructural or structurally open. They are *hypernonmonotonic*. We should understand that material reason relations in general can include hypernonmonotonic implications.

It might well be sensible to require the inconsequentiality of explicitation as a structural constraint on *logical* consequence relations. But just as for the logical expressivist there is no good reason to restrict the rational relations of implication and incompatibility we seek to express with logical vocabulary to monotonic or transitive ones, there is no good reason to restrict our expressive ambitions to consequence relations for which explicitation is inconsequential. On the contrary, there is every reason to want to use the expressive tools of logical vocabulary to investigate cases where explicitation *does* make a difference to what is implied.

One such case of general interest is where the explicit contents of a premise-set are the records in a *database*, whose implicit contents consist of whatever consequences can be extracted from those records by applying an *inference engine* to them. (The fact that the "sentences" in the database whose material consequences are extracted by the inference engine are construed to begin with as *logically* atomic does not preclude the records having the "internal" structure of the arbitrarily complex datatypes manipulated by any object-oriented programming language.) It is by no means obvious that one is obliged to treat the results of applying the inference-engine as having exactly the same epistemic status as actual entries in the database. A related case is where the elements of the premise-sets consist of experimental *data*, perhaps measurements, or observations, whose implicit content consists of the consequences that can be extracted from them by applying a *theory*. In such a case, explicitation is far from

inconsequential. On the contrary, when the CERN supercollider produces observational measurements that confirm what hitherto had been purely theoretical predictions extracted from previous data, the transformation of rational status from *mere* prediction *implicit* in prior data to actual explicit empirical observation is an event of the first significance—no less important than the observation of something incompatible with the predictions extracted by theory from prior data. This is the very nature of empirical *confirmation* of theories.

Imposing Cut and Cautious Monotonicity as global structural constraints on material consequence relations amounts to equating the epistemic status of premises and conclusions of good implications. But in many cases, we want to acknowledge a distinction, assigning a lesser status to the products of risky, defeasible inference. In an ideal case, perhaps this distinction shrinks to nothing. But we also want to be able to reason in situations where it is important to keep track of the difference in status between what we take ourselves to know and the shakier products of our theoretical reasoning from those premises. This idea is one important motivation for Strict-Tolerant logics, paradigmatically, ST (Cobreros et al., 2013). Chapter Five discusses how this logic shows up from the vantage point of the more expressively powerful treatment of conceptual roles made possible by our implication-space semantic metavocabulary. Our expressive aspiration is for a logic that can codify reason relations that do not presuppose or build in a structural requirement of explicitation closure: the demand that explicitation always be inconsequential.

2.6 Explicitation Closure and Rational Hysteresis

Let us take stock. We began by considering two diametrically opposed approaches to understanding the relations between logic and reasoning. Taking on board the idea that logic concerns the reason relations of implication and incompatibility that govern reasoning practices and processes sharpened the issue somewhat. Logicism claims that logic determines the proper relations of implication and incompatibility: that implications hold rationally just in case they are or are supported by *logically* good implications (valid deductions), and that rational incompatibilities are or are supported by *logical* incompatibilities, that is, formal inconsistencies. Expressivism understands the task of logic to be expressing material, prelogical reason relations of implication and incompatibility: making them explicit in the sense of sayable, claimable contents, for which reasons can be asked and given. The rational importance of logical vocabulary is not that it lets us prove theorems of logic that determine what relations of implication and incompatibility really hold, but that the sentences of any logically enriched prelogical base

vocabulary let us say what relations of implication and incompatibility hold in that prelogical base vocabulary.

We then considered a weaker form of logicism: *structural* logicism. This is the view that the algebraic structure of material reason relations is the same as the algebraic structure of specifically logical reason relations: implications and incompatibilities that hold just in virtue of the logical form of the sentences involved. For the case of logical consequence, it is generally agreed that it has the structure of a topological closure operation. Combining Tarski's and Gentzen's versions of these structural principles (so as to extrude irrelevant details particular to their formulations), this means that consequence relations satisfy Containment (CO), Monotonicity (MO), and Cumulative Transitivity (CT). The traditional form of structural logicism accordingly takes implication to have a *topological closure* structure. Assuming Tarski and Gentzen are right about the structure of logical consequence (and this much is not at issue, for instance, between classicists and intuitionists), logicism about implication generally entails the topological closure form of structural logicism. If structural logicism is wrong about implication in general, if material consequence relations do not in general exhibit the topological closure that comprises CO, MO, and CT, then logicism about reasons in general cannot be right either.

We then argued that, however it might be with specifically *logical* implication, *material* consequence relations are not in general monotonic. That is enough to show that the topological closure version of structural logicism is not true. But what structure do consequence relations in general exhibit? It was pointed out that a popular principled fallback from MO is Cautious Monotonicity (CM).[13] It is the principle that no consequences of a premise-set Γ are lost by adding any collateral premises that are already consequences of Γ. CM seems a particularly natural weakening of MO to consider, because it is dual to CT. CT says that adding consequences of Γ to Γ never *adds* any new consequences, and CM says that adding consequences of Γ to Γ never *subtracts* any consequences. Observing this duality brings into view the operation they share: moving a sentence across the implication turnstile, from being a conclusion to being a premise. Thinking of what a premise-set implies as what it contains *implicitly* (its implicit content) and the actual elements of the premise-set as what it contains *explicitly* (its explicit content) makes it natural to call this process "explicitation." It is making implicit content explicit: a prelogical sort of expression.

Explicitation, in turn, makes visible a further kind of closure structure: explicitation closure. If both CM and CT hold globally for an implication relation, then explicitation is guaranteed to be *inconsequential*. Making implicit content explicit never affects implicit content, neither increasing or decreasing it. This sort of closure is weaker than topological closure,

just insofar as CM is weaker than MO. But it, too, expresses commitment to a kind of *stability* of consequences, a kind of closure.

And with this weaker sort of closure structure comes a new sort of structural logicism: *explicitation-closure* structural logicism (now distinguished from *topological*-closure structural logicism). This is the claim that consequence relations in general have at least this much of the structure of logical consequence: CM and CT, as well as CO.[14] Explicitation closure might indeed seem to be an attractive fallback from topological closure as a candidate for being the structure of consequence relations generally. It is implied by the full topological closure Tarski and Gentzen take to be characteristic of logical consequence, and would permit acknowledgment of a structural distinction between logical and material consequence relations, if the material ones satisfy only the weaker explicitation-closure structure, rather than the full topological-closure structure of logical consequence as traditionally conceived.

To argue against the explicitation-closure form of structural logicism we introduced a special case of consequential reason relations to serve as a model. It can be helpful to think of the explicit premises of an implication as a database, and the turnstile as standing for a theory functioning as an inference-engine that, when applied to the database, yields the conclusions, thereby extracting the content implicit in the database. Thinking of instances where the database contains observational data and the inference-engine extracts the predictions of some theory shows that the inconsequentiality of explicitation is *not* a condition we want to insist on for implication or consequence relations in general. The difference in status between what has been observed or measured and what is merely predicted by theory is too important to have the boundary between them structurally erased. Explicitation is not in general or *de jure* inconsequential: it can make a difference to what is implied by a premise-set. The explicitation-closure conditions CM and CT together imply that explicitation is inconsequential in this sense. So this weaker form of closure should also be rejected as a global structural constraint on consequence relations, and explicitation-closure structural logicism must accordingly be rejected along with topological-closure structural logicism.

Remember the issue that led to addressing the issue of whether material consequence relations should be understood to be structurally closed in the first place (whether in the topological or the explicitation sense). It was to understand the constraints on *expressing* material reason relations using *logical* vocabulary. The database + inference engine model reminds us that we should aim for the greatest possible flexibility and expressive power possible. We should build in as few *a priori* constraints on logically codifiable inference engines as possible. The expressivist's aim should be to produce and deploy logical tools for expressing reason relations of all

intelligible structures. The expressivist ideal is to develop the expressive power to make explicit *any* and *all* species of the turnstile, any and all senses of "follows from." Thinking in terms of databases and inference engines reminds us of just how capacious that class (and so that aspiration) is. Perhaps it is a utopian aspiration. (Spoiler: it is not.) From this point of view, CM should not be assumed to hold globally, any more than CT should. From an expressivist point of view, we want a logic that can be introduced conservatively over, and has the expressive power to codify, material reason relations that are *open*, not just in not being *topologically* closed, but in not being *explicitation* closed, either. Structural logicism in both forms must be rejected.

The material consequence relations that the expressivist takes it logic should aim to codify are accordingly radically substructural, in that it should not be presupposed that material, nonlogical relations of implication satisfy global structural principles of the sort characteristic of specifically *logical* consequence relations. The kind of structure denied is *closure* structure, of the two sorts distinguished here: topological closure and explicitation closure. The substructural consequence relations they contrast with are *open*, in both corresponding contrasting senses. The significance of reason relations with open structure is best grasped by thinking about the process of explicitation. Explicitation as a process is *drawing* or extracting consequences, and adding them to the premises from which one reasons. This is a central, indeed, essential kind of *inference*: acknowledging *explicitly* (by treating them as premises in further implications) conclusions that were *implicit* in some prior premise-set—"implicit in" in the literal sense of "implied by." (Harman reminds us that this is not the only rational inferential process governed by implication relations. Denying what is in this sense implicit in a premise-set can also entitle one explicitly to deny some or all of the premises.) The key point is that in an open structural setting, making explicit any set of consequences of a premise-set might add some new consequences (where there are local failures of CT) and subtract others (where there are local failures of CM), relative to the consequences of the original premise-set. This has a number of consequences, which highlight the striking differences between the general case and structurally closed, specifically *logical* consequence relations.

Most importantly, realizing that explicitation is not in general inconsequential, that failures both of CM and of CT are to be expected, challenges the common philosophical conception of the *rational* closure of a knower's beliefs. It has seemed natural to some to think that rationality requires that one believe all the consequences of one's beliefs or commitments: everything that follows from them. In a familiar dialectic, opponents point out that this is an impractical demand. There are far too

many such consequences, and simply surveying them would take too long. In response, while it is acknowledged that it might be impossible actually to achieve the rational closure of one's beliefs or commitments (attitudes of accepting or denying claimables), so that treating doing so as constitutive of rationality would entail that no-one ever is or could count as fully rational, nonetheless it is suggested that such closure could still be understood as a regulative ideal of rationality. It is then objected (for instance, by Harman) that we should not understand reasoning, in the sense of drawing further consequences from our beliefs, as governed by a norm we can never in practice live up to. The important thing to realize about this debate is that it assumes that the concept of all the consequences of a premise-set (its rational closure in the sense of everything that follows from it) makes sense or has a definite reference. The argument has been about whether closing one's beliefs under rational consequence is practical, desirable, a regulative ideal of rationality, and so on. But if CM or CT sometimes fail, then there is in general no such thing as *the* set of consequences of a premise-set. The dialectic only makes sense assuming at least the weaker form of structural logicism: that, as for purely logical implication, explicitation is in general inconsequential.

If material reason relations can be hypernonmonotonic and nontransitive, then rational consequence does not have a closure structure. There need be no closed set of consequences for each premise-set, which cannot be added to or subtracted from by continuing to extract consequences. And that means that the process of explicitation—of explicitly acknowledging consequences of the claimables one accepts, and using them as premises in the drawing of further consequences for which they provide reasons—becomes an open-ended enterprise. There is no unique endpoint that one is guaranteed eventually to reach, by just keeping at the process of explicitating rational consequences indefatigably, and making no mistakes.

- Acknowledging some consequences one's commitments give one reasons for can put one in a position where one no longer has sufficient reasons to draw other consequences that one could otherwise have drawn (that one would otherwise have been entitled to draw) from the original set.

That is a consequence of the failure of CM to hold globally.

- And explicitating to begin with only some of the consequences of the original premise-set might well provide reasons for new consequences, that the original premises, taken all together, did not imply or directly provide reasons for.

That is a consequence of the failure of CT to hold globally.

If explicitation closure is not guaranteed by the structure of implication as such, then there is no guarantee that the process of explicitating consequences, and then explicitating the consequences of *that* expanded premise-set, and then explicitating the consequences of *that* set, and so on, will reach a fixed conclusion. In open settings one cannot be sure in advance that the process of explicitation will reach a stable stopping-place—that it will arrive at a premise-set all of whose explicitations are inconsequential, involving no violations of CT or CM. It can happen that every position $X \cup Y$ one arrives at by explicitating consequences of the original premise-set X (CO guarantees one will arrive only at supersets of X) still has some implicit content, some set of consequences, such that when they are added to $X \cup Y$ as explicit premises, the result is a different set of consequences than $X \cup Y$ had. The process of explicitation need not end. The explicitation closure conditions, by contrast (like the topological closure conditions), guarantee *finality*: that $Cn(X) = Cn(Cn(X))$ for every premise-set.

A further observation is that with open consequence relations there is nothing privileged about the result of explicitating *all* of X's consequences at once. One might get a larger implicit content by explicitating only some of the consequences of a premise-set. Closed consequence relations ensure that the result of explicitating all the consequences of a premise-set will include the result of explicitaing only some of them. So not only is finality guaranteed, but one can get to the final fixed point of the explicitation process in a single step. Explicitation of topologically closed consequence relations is *immediate*.

Since finality is not guaranteed, explicitation of consequence relations with open structure can be radically path-dependent. Nor is there any guarantee that following any particular explicitation path that starts by adding those consequences will ever arrive at any of the premise-sets or consequence sets reached by starting with the explicitation of a different proper subset of the consequences of Γ. As a result, the consequences that come into view at any point in the process of explicitating some premise-set Γ depend on which of Γ's consequences one chooses to explicitate first. This is explicative *hysteresis*. By contrast, explicitation of structurally closed consequence relations is *stable*, in the sense of being path-independent. Since it is *final* there is a fixed endpoint to the process: $Cn(X)$. Since it is *immediate* one can jump to that endpoint in a single explicitating step. And since it is *stable*, if one *did* explicitate step-wise, the results would be the same endpoint, for all explicitation paths from X to $Cn(X)$.

Finality, immediacy, and stability of explicitation are all very useful features of consequence relations. There are many good reasons to want to build them into the *logical* consequence relations that govern the use of the logical vocabulary whose expressive task it is to make explicit the reason

relations that govern the use of all kinds of vocabulary. (The purely logical consequence relation determined by the rules of our favored logic NMMS does have all these nice closure properties—though the full consequence relation of the logical extensions of arbitrary base vocabularies does not.) But if, through a thoughtless and misplaced commitment to either variety of structural logicism, we project those ideals onto the actual material consequence relations that govern the use of nonlogical vocabulary and articulate the conceptual contents they express, we make invisible essential structural features of the crucial rational *process* of *drawing* consequences from premise-sets, of acknowledging explicitly the implicit, consequential contents of explicit commitments. The process of explicitation is an important species of inference. Studying it is one of the reasons we want logical tools to make reason relations explicit. The most elementary sort of pragmatism counsels that we do not obscure rational processes and practices by imposing Procrustean *a priori* structural restrictions derived from mistakenly taking specifically *logical* reason relations of implication and incompatibility as the model on which to understand the structure of reason relations in general.

The fundamental structural feature of the process of explicitation that is being highlighted here is that sequentially extracting consequences from a premise-set—explicitly acknowledging them, and then extracting more of the implicit content of that set—is not a process that is guaranteed either to converge or, in the cases where it *does* converge, to reach a *unique* final set of consequences.

Reasoning in this explicative sense of making explicit what is rationally implicit in a set of commitments is in principle severely *path-dependent*. Where one ends up depends upon what choices one makes about what to explicitate *first*, and in general on the *sequence* of explicitations one effects. Indeed, those choices determine even whether an inferential process of extracting and acknowledging consequences converges at all, as well as where it ends up. The historical path-dependence of consequences and incompatibilities is a deep *structural* feature of material, nonlogical reasons as such.

In the physical sciences, the technical term for path-dependence of processes is "hysteresis." That is why the phenomenon in question here is called the *rational hysteresis* of explicitation—the hysteresis of making explicit or acknowledging explicitly what is rationally implicit in (implied by) a set of commitments. It ensures that *history* is an intrinsic feature of *reason*. This is the *rational* (not yet logical) basis of the *historicity of reason*. It is an essential element of the explanation of why *discursive* creatures, creatures who can *talk*—that is can engage in practices of making claims and challenging and defending them—have *histories*, and not just *natures*.

It was pointed out above that material incompatibility relations can be nonmonotonic, just as material consequence relations can be. Material incompatibility relations can also be hypernonmonotonic. Base vocabularies can have premise-sets for which not only MO fails, but Cautious Monotonicity fails. These are sets of sentences whose status as coherent or incoherent can be changed by explicitation of some or all of their consequences. Perhaps the most interesting case is coherent sets that are implicitly incoherent, in that adding one or more of their consequences as explicit premises results in an explicitly incoherent set. Many, perhaps most, philosophy books are probably like this. Their explicit content is expressed by the sentences that occur on their pages. Their implicit content is what those sentences jointly imply but do not explicitly state. Part of the work of readers, reviewers, and seminars studying the text is to extract at least some of its consequences and critically confront the explicit text with challenges presented by incompatibilities between the explicit content and what it implies. Charitable interpreters will consider whether the incoherences that emerge from such exercises are curable by the addition of further premises.

For the implicit incoherence of a premise-set need not be *persistent*. It might be curable or defeasible by the addition of further premises. A standard strategy for curing an incompatibility by adding a further premise is epitomized in the Scholastic maxim "When faced with a contradiction, make a distinction." (That is roughly how functionalists in the philosophy of mind proposed to reconcile the materialist's claim that thoughts must be specifiable in the language of physics or neurophysiology with the dualist's claim that what makes something the thought that *p* cannot be specified in the language of physics or neurophysiology.) The possibility under consideration is that the distinction required to reconcile incompatible claims might be implicit in those claims themselves. (The claim that, at least within some domains, *every* contradiction implicitly contains a distinction that would resolve it is a recognizably Hegelian one.) Some additional premises that would, if explicitly added to the premise-set, result in a coherent consequence set might sometimes be found even among those implied by the incoherent result of prior explicitation. Exploring the many paths in and out of the dangerous territory of explicit incoherence that explicitation of implicit content can trace out is an interpretive adventure. It is characterized by what we are in a position to recognize as hermeneutic hysteresis (not to be confused with hermeneutic hysteria). Among the way-stations along many such explicitation paths will be found explicitly incoherent sets of claims that are implicitly coherent. In such cases, the incoherence is corrigible by adding premises that are already implied by an explicitly incoherent premise-set. The difficulties of the text as it stands can be rectified by merely explicitly drawing consequences to which the text

already leads. The best readers, editors, and referees will be most helpful when pointing these out.

Exploring the material reason relations that articulate the conceptual contents expressed by the use of nonlogical vocabularies by traversing explicative pathways leading from one premise-set to another is a principal form of semantic understanding—understood now as a process rather than a state. Expressivists want a logic expressively powerful enough to make explicit what we are doing and what we find out when engaged in this sort of hermeneutic enterprise. That requires distinguishing at least the three kinds of incoherence/incompatibility just pointed out: explicit, implicit, and persistent. The very existence of these important phenomena is rendered invisible by structural logicism. Traditional logics are inadequate to the task of expressing these hypernonmonotonic, open-structured regions of material reason relations.[15]

2.7 Conclusion: Expressivist Criteria of Adequacy for Logical Vocabulary

This chapter began by contrasting two different ways of thinking about the relations between logic and reasons in general. Logicism about reasons understands "good reason" to mean "logically good reason." Reasons for commitments are governed by logically valid implications, and reasons against commitments are governed by logical inconsistencies. Where it is not obvious on the surface that this is so for some nonlogical reasons, if reasons are to be understood to be involved in the use of that nonlogical vocabulary, a logical deep structure must be discerned as underlying its use. By contrast, expressivism about logic treats the material goodness of prelogical reason relations of implication and incompatibility as prior in the order of explanation to reason relations of specifically *logical* consequence and inconsistency. The defining task of logical vocabulary is an expressive one: to make explicit, in sentential, claimable form, the reason relations of material consequence and incompatibility that normatively govern the use of sentences expressing claimables in nonlogical base vocabularies. Logical reason relations are *elaborated from* the reason relations of base vocabularies, and the logically complex sentences whose use those logical reason relations govern are *explicative of* the implications and incompatibilities of those base vocabularies. We say that logical vocabulary is "LX" for the base vocabularies from which it is *el*aborated ("L") and of which it is *ex*plicative ("X"). An expressively ideal logical vocabulary would be *universally* LX: LX for *any* base vocabulary (lexicon of sentences governed by reason relations) whatsoever.

The bulk of this chapter has been concerned with a structural stumbling block that confronts both the logicist and the expressivist programs.

The algebraic structure of specifically logical reason relations is quite different from the algebraic structure of reason relations of implication and incompatibility in general. Logical consequence has a topological closure structure. It is monotonic, transitive, and reflexive (or satisfies Containment, sometimes called "Contexted Reflexivity"). Nonlogical implications, by contrast, can be nonmonotonic and nontransitive. Indeed, they can be hypernonmonotonic, in the sense that the implication of a conclusion by a premise-set can be defeated or infirmed by making some of its implicit content explicit: adding as an explicit premise something that was already implied by the original premise-set. Material incompatibilities can also be nonmonotonic, and even hypernonmonotonic. The logicist about reasons is obliged somehow to construct substructural (open-structured) reason relations out of the structurally closed relations of logical consequence and inconsistency. That is the task of traditional nonmonotonic logics.

This structural mismatch between material and logical reason relations also poses challenges for both sides of the expressivist ideal for logical vocabulary: that it should be capable of being elaborated from and explicative of the reason relations of every material base vocabulary. On the side of elaboration, if logically complex sentences using conditionals and negations are to be capable of expressing explicitly material relations of implication and incompatibility, their addition must conservatively extend the reason relations of the base vocabulary. If adding the new logical vocabulary changed the material reason relations of the base vocabulary, for instance by imposing further structure on them, that would undermine the capacity of the new logical vocabulary to express them explicitly. But some standard connective rules force the consequence relation governing the whole logically extended vocabulary to be monotonic, in the context of other plausible principles. For example, part of Gentzen's rule for introducing conjunction on the left side of sequents (put in the notation used here) is:

$$\frac{\Gamma, A \mid\!\sim \Delta}{\Gamma, A \wedge B \mid\!\sim \Delta}$$

If $\Gamma, A \wedge B \mid\!\sim \Delta$ entails $\Gamma, A, B \mid\!\sim \Delta$ (as seems plausible), then this permits a transition from $\Gamma, A \mid\!\sim \Delta$ to $\Gamma, A, B \mid\!\sim \Delta$ via the intermediate step of $\Gamma, A \wedge B \mid\!\sim \Delta$, even if all of Γ, Δ, A, and B are sentences of the nonlogical base vocabulary. This just is monotonicity. In this case, one could avoid building in monotonicity into the connective definition by using a "multiplicative" rule rather than the "additive" one:

$$\frac{\Gamma, A, B \mid\!\sim \Delta}{\Gamma, A \wedge B \mid\!\sim \Delta}$$

If one imposes monotonicity as a global structural constraint, the additive and the multiplicative rules are equivalent. If not, they come apart—and the difference becomes important. The point remains: expressivist logics must be able to accommodate base vocabularies that have radically open structures: nonmonotonic, hypernonmonotonic, and nontransitive, without imposing additional structure on the original reason relations. The structural disparity of logical and material reason relations makes the conservativeness condition on the elaboration of logical reason relations from material base reason relations challenging to satisfy.

On the side of explication, expressivist logics aim to be capable of explicitly expressing (putting in statable, claimable form) the radically open-structured reason relations of any and all base vocabularies, even where they are hypernonmonotonic, nontransitive, and fail to satisfy other minimal global structural conditions such as Containment. That principle says that all sequents of the form $\Gamma, A \mathrel{|\!\sim} A, \Delta$ are good: that what a premise-set contains explicitly, it also implies, and so contains implicitly. That is a plausible, because severely restricted, monotonicity principle, allowing the weakening of instances of Reflexivity of the form $A \mathrel{|\!\sim} A$ by arbitrary Γ and Δ. But it is rejected by relevance logicians (due precisely to the irrelevance of Γ and Δ to the goodness of the implication), and expressivists aspire to be able to codify the reason relations of relevance logics, just like any other base vocabulary. Sufficiently expressively powerful logics should even be able to handle failures of Gentzen's principle of Contraction:

$$\frac{\Gamma, A, A \mathrel{|\!\sim} \Delta}{\Gamma, A \mathrel{|\!\sim} \Delta}$$

That principle is not satisfied by the consequence relation of Linear Logic. (Its inventor, Jean-Yves Girard says that anyone who accepts Contraction deserves "two kicks in the ass—not two occurrences of the same kick" (Girard, 2001, 482).) Expressivists want logics that can make explicit the reason relations of Linear Logic, just as for any other base vocabulary.

This is a tall order. Chapters Three and Five show how this expressivist aspiration can be fulfilled, nonetheless. Astonishingly, familiar logics can be tweaked so as to be demonstrably LX for all these sorts of radically substructural reason relations—and are supraclassical (for the multisuccedent case) or supraintuitionist (for the single-succedent case)—for base vocabularies that satisfy Containment. Further, the implications and incompatibilities that hold in virtue of logic alone, in the sense of holding no matter what base vocabulary the logically extended vocabulary is elaborated from, can be fully structurally closed: both monotonic and transitive. Of course, there are many implications and incompatibilities

that hold in the logically extended language that *do* depend on the base vocabulary, and do not hold in virtue of the logical connective definitions alone. It is these we are particularly interested in. For it is these nonlogical reason relations holding in the logically extended vocabulary that make it possible to *say* there what the reason relations of the base vocabulary (and also its logical extension) are.

In addition to making explicit the material consequence and incompatibility relations that govern the use of arbitrary open-structured or substructural base vocabularies, expressivist logics also aspire to mark explicitly any local regions of those reason relations that do exhibit structural features that do not hold globally. Even material consequence and incompatibility relations that are not globally monotonic or transitive typically include some particular implications and incompatibilities that are persistent, that is that do hold monotonically. Some implications will satisfy Cumulative Transitivity. The regions of the reason relations where both monotonicity and transitivity hold amount to local islands of topological closure of reason relations within the larger sea of open-structured ones. Structural logicism about reason relations has obscured the intricate open structure of the implications and incompatibility that actually articulate the conceptual content of much material, nonlogical vocabulary. Expressing these features of the fine structure of material reason relations of base vocabularies in extensions of those vocabularies elaborated by introducing logical vocabulary is also something expressivists would like their logics to be able to do. Chapter Three shows how this expressivist criterion of adequacy can also be satisfied. Chapter Five introduces implication-space semantics, which is a model-theoretic metavocabulary with the expressive power to make explicit the fine structure of the conceptual roles that sentences of open-structured base vocabularies play in virtue of standing in reason relations of these intricate kinds.

Mentioning these two dimensions along which expressivist logics can hope to make reason relations explicit—saying *that* particular implications and incompatibilities hold and explicitly marking implications and incompatibilities that exhibit structural features that hold only locally and not globally within particular constellations of material reason relations—points in the direction of a response to a worry that residual structural logicist commitments and temptations might raise about the very idea of open-structured reason relations. For Tarski and Gentzen, and most of their heirs and admirers, the topological closure structure of logical consequence relations is essential to anything that deserves to be called a *consequence* relation. Similar remarks apply to the monotonicity of inconsistency, and even the explosion principle *ex contradictione quodlibet* concerning the interaction of implication and incompatibility relations. If,

as we are recommending here, we drop *all* of those structural constraints on reason relations, what remains of the claim that they are *reason* relations? It might seem that we are just talking about relations, *tout court*. We are insisting that there are two kinds of reason relations, what we *call* "implication" and "incompatibility." The latter are symmetric, and the former are not. But what is it about a set of relations, some symmetric and some not, in virtue of which they deserve to count as "reason" relations? Why isn't this a case of throwing the baby of reasons out with the bathwater of closure structure? If structural logicists are wrong in appealing to structural features such as monotonicity and transitivity as essential features demarcating *rational* relations, what *does* pick out some relations among sentences as rational relations—distinguishing them from, say, orthographic or syntactic relations?

We have been suggesting that what is of most philosophical interest, and what logical expressivists should be most interested in being able explicitly to codify in the form of sentences that can themselves stand in reason relations, is the fine structure of the material reason relations exhibited by particular nonlogical vocabularies, such as regions of subjunctive robustness (local persistence of implication or incompatibility) and behavior under explicitation of implicit consequences. These intricate, detailed rational relations among sets of sentences only become visible once the traditional global structural demands that amount to requiring topological closure are relaxed. But if all that remains at the level of grosser global algebraic structure is so thin and abstract, doesn't that undercut the claim that what we are talking about is still relations of implication and incompatibility? What does it mean for the relations we are discussing to *be* relations of implication and incompatibility, what we have been calling "reason relations"?

We take this challenge very seriously. One of the principal tasks of this book is the development of an extended response to it. As we say in the Introduction, the form of that response is a kind of functionalism about reasons at the metatheoretic level. What makes the open-structured relations we discuss qualify or deserve to count as *reason* relations of implication and incompatibility is the role they play with respect to four kinds of metalanguage. We saw in Chapter One how relations of implication and incompatibility can be understood in terms of the use of sentences to make claims and defend and challenge them, when those discursive practices are specified in a bilateral normative *pragmatic* metavocabulary. That account neither entails nor presupposes that those relations have an algebraic or topological closure structure. Chapter Three shows how even the radically open-structured relations we are calling "implication" and "incompatibility" can be codified in a *logical*

metavocabulary in the form of a familiar sort of multisuccedent sequent calculus that is elaborated from and explicative of arbitrary material base vocabularies, regardless of whether their structure is open or closed in the various senses we have distinguished. Chapter Four shows how to understand implication and incompatibility in a Finean truth-maker *semantic* metavocabulary from which all requirements of closure structure have been removed. It further reveals a remarkable isomorphism between that truth-maker semantic specification of reason relations and the bilateral normative pragmatic specification of them. Chapter Five then offers a metavocabulary for defining and manipulating the *conceptual roles* (showing up there as "implicational roles") sentences can play in virtue of the open-structured reason relations they stand in to one another. Our *metarational functionalist* claim is that a good thing to mean by "reason relations" is whatever plays all four of these roles: pragmatic, logical, semantic, and in the specification of conceptual roles, regardless of how spare and abstract the gross global structural constraints imposed on those relations might be. If it walks like a duck (pragmatically), quacks like a duck (in expressive logics), looks like a duck (semantically), and lays eggs (conceptual roles) like a duck, one is entitled to call it a duck.

This metalevel functionalism about reason relations also highlights a substantial difference between the very abstract account of reasons presented here and substantive contemporary accounts of rationality such as Bayesianism and rational choice theory. Given the prominence and popularity of such programs, our focus here on logicism as a way of thinking about reasons, and so rationality, might have seemed odd, blinkered, or perhaps merely old-fashioned. After all, for many contemporary thinkers, classical logic only matters insofar as it is a limiting case of probability theory, resulting from allowing probabilities to take only the values 1 and 0. These theories indeed formally define clear and precise ways of understanding the concept of good reason. But they are engaged in a very different sort of enterprise than the one pursued here. This is clear from the fact that the sentences to which credences are attached are taken to be already conceptually contentful, in advance of playing the role in reasoning Bayesians reconstruct. In rational choice theory, the sentences specifying options and outcomes must already be semantically interpreted in order to support assignments of conditional probability of outcomes relative to options, and preference and utility orderings on outcomes. The sense of "good reason" these theories define is accordingly not a candidate for explaining the semantic contentfulness of the sentences whose role in good reasoning is being defined. The semantics is outsourced, or rather just presupposed as available when the Bayesian or rational choice theorist comes on the scene. In this sense, our spare, highly abstract conception of reason relations aims to be explanatorily more basic than and conceptually

prior to the late-coming substantive conceptions of rationality and good reason these approaches articulate.

This orienting aspiration—to specify a conception of reason relations that is capable of underwriting an understanding of the conceptual contents of sentences as roles they play in reason relations—might well motivate a skepticism about our enterprise based on the liberality of our abstract specification of reason relations, rather than skepticism based on commitment to structural logicism. For we are thinking to begin with of what stands in reason relations as just sentences as lexical items (what Sellars calls "sign designs"), not as the conceptual contents or propositions they express. How, one might wonder, can relations among such initially meaningless items (corresponding to Wittgenstein's "sign-post considered just as a piece of wood") be intelligible as *reason* relations of implication and incompatibility? What could justify calling them that? The programmatic commitment in question is an essential feature of the pragmatics-first order of explanation that seeks to explain how discursive practices of using sentences to make claims and challenge and defend them by making further claims that thereby have the practical significance for other practitioners of serving as reasons for or against them can confer conceptual contents on those sentences. We defend it by specifying the multidimensional metatheoretic role we claim characterizes reason relations as such. We have already offered a pragmatic story about the role in such practices relations among sentences must play in order to qualify as reason relations of implication and incompatibility. It will be supplemented in the chapters to come by an account of how reason relations in that pragmatic sense can be expressed explicitly in logical metavocabularies, using conditionals to express implications and negation (with conditionals) to express incompatibilities, as well as using modal operators to express local regions of structure. Abstract, radically substructural material reason relations then show up from a different metatheoretical perspective as interpretable in robust model-theoretic truth-maker semantic metavocabularies that can be proven to pick out the very same relations of implication and incompatibility previously picked out pragmatically and expressed logically. Formally tractable and usefully manipulable conceptual roles fit for substantial further explanatory work turn out to be definable from those very same abstractly characterized open-structured relations. In the next three chapters, in the course of elaborating specially constructed versions of these different kinds of discursive metavocabulary, we demonstrate a web of systematic relations among them that together define a precise sense in which they provide perspectives on a common topic: reason relations of implication and incompatibility.

Notes

1 In this work we do not consider *practical* reasons, or reasons for commitments other than specifically doxastic ones.
2 This idiom is introduced in Brandom (2008), Chapter Two.
3 See Belnap (1962). Complications to this analysis ensue in substructural settings—in particular, nontransitive ones, where Cut fails. See Ripley (2015).
4 This line of thought is presented in more detail in Chapter One of Brandom (2000), and Chapter Two of Brandom (1994).
5 From our point of view, "negation on the Australian plan" is just logical expressivism considering only one bit of logical vocabulary: negation.
6 Tarski put forward a model-theoretic version of this approach that appeals to permutation rather than substitution, which has been further developed and championed in our own time by Gila Sher (1991). At a very abstract level, the difference between these is whether we look at variations *salva consequentia* or *salva veritate*.
7 https://protect-us.mimecast.com/s/ZfbuCADQm5Ur4XPZjU9Ttvg?domain=aeon.co
8 As explained below, their projects of building nonmonotonic logics out of classical monotonic ones is alien to the expressivist approach. Expressivists want logics that are expressively adequate to codify nonmonotonic reason relations of implication and incompatibility. That is not at all the same enterprise as constructing *logics* that are nonmonotonic. In fact it turns out that the purely logical reason relations governing such logics (the implications and incompatibilities that hold in virtue of logic alone) governing logics that codify nonmonotonic material consequence relations can themselves be monotonic and transitive: structurally closed. So in the end we are in a position to justify the invocation of classical negation and inconsistency as part of the definition of one's nonmonotonic logic—if one still wants to do what these nonmonotonic logicians want to do.
9 We discuss the codification in implication-space semantics of reason relations of this sort near the end of Chapter Five, in Section 5.5.4.
10 This case has been made forcefully by Dov Gabbay (1985), who includes also CO and CT as necessary for workable nonmonotonic systems.
11 "A version of" because CT is additive, that is, context sharing, while Gentzen's Cut was multiplicative, that is, context combining. In open, substructural settings, these diverge.
12 This is a distinct, new notion of explicitation, not to be confused with the one that is involved in LX-ness. It does not specifically involve logical vocabulary. In Chapters Three and Five we will introduce still further layered senses to this multifarious concept, which is at the center of our expressivism.
13 We do not discuss the other most popular candidate weaker than MO, often called "rational monotony," because as usually formulated, it assumes the language already has negation in it. It says that no conclusions of a premise-set are lost by adding new premises that do not contradict it. There is a version that appeals only to incompatibility, but looking at structural principles relating

implication and incompatibility would take us too far beyond the argument here.

14 CO comes from Tarski's plausible condition that $X \subseteq Cn(X)$: any premise-set is contained in its consequence set. In the context of MO, CO is equivalent to Reflexivity (RE), which is what Gentzen actually uses (all leaves of all purely logical sequent derivations are instances of RE). If we relax MO, this equivalence breaks down, and one might worry, as relevance logicians do, about endorsing even the weak sort of monotonicity that CO enforces: monotonicity of implications of the form $A \mid\!\sim A$, which can be arbitrarily weakened with further premises. But CO remains plausible when thought of in explicitation terms: what a premise-set *explicitly* contains counts as also *implicitly* contained in it. Explicit content is part of implicit content.

15 To jump ahead in our story a bit: the Incoherence-Incompatibility (II) codifying definition of negation introduced earlier entails that an incoherent set implies the negations of all of its members. If we explicitate those implications, add any of those negations as explicit further premises, the result is a premise-set that for some sentences contains both A and $\neg A$. Such premise-sets are *persistently* explicitly incoherent. If implicational explosion is retained for persistently incoherent sets, then taking the explicative path through these logically complex negated sentences will close off the possibility of further exploration by explicitation. The capacity to make reason relations explicit is a two-edged sword. With great expressive power comes great hermeneutic responsibility.

3 Introducing Logical Vocabulary
Making Reason Relations Logical

Ulf Hlobil

In this chapter, we introduce a broadly proof-theoretic way to codify and theorize open reason relations among sentences of a language that contains logical vocabulary. We interpret this proof-theoretic formalism in a pragmatic-normative way. Throughout this book, we restrict ourselves to the sentential level, as things turn out to be fascinating and complex enough at that level. Thus, we are mainly concerned with negation ($\neg$), the conditional ($\rightarrow$), disjunction ($\vee$), and conjunction ($\wedge$). In this chapter, we present a sequent calculus treatment of these connectives that is almost entirely standard and conventional. What is exciting about this treatment, however, is that it allows us to capture, in a formally rigorous way, the ideas of logical expressivism that we have presented in the previous chapters, and this includes capturing open reason relations in which topological closure as well as explicitation closure fail. In later chapters, we will often return to the logical systems that we introduce in this chapter. One of the chief accomplishments of this book lies in the fact that we show how these logical systems appear and emerge naturally from different perspectives, including the perspective of truth-maker theory in the next chapter and the more abstract perspective of implication-space semantics in Chapter Five.

Before turning to more formal and technical matters, let us recapitulate where we are in our philosophical project. In the first two chapters, we have introduced the idea of reason relations according to a pragmatics-first approach, as well as the thesis of logical expressivism. To repeat, we start with a discursive practice of giving reasons for and against claims. This practice is constrained by reason relations of implication and incompatibility. In particular, the sentences in a set Γ are jointly a reason for the sentence A if and only if Γ implies A, which we write as $\Gamma \mid\!\sim A$. And the sentences in Γ are jointly a reason against the sentence A if and only if Γ is incompatible with A, which we write as $\Gamma \# A$. We encode both relations in one single relation by stipulating that $\Gamma, A \mid\!\sim$ (notice the empty right side)

just in case $\Gamma \# A$. According to logical expressivism, it is the characteristic function of logical vocabulary to make such reason relations explicit.

The sense of "reason for" and "reason against" in this book is emphatically not the sense in which a reason for or against a claim makes it the case that one ought to accept or reject the claim, respectively. For, what one ought to do may be to reject the reason, as Harman (1986) has stressed in his work on the normativity of logic. Rather, if one accepts or rejects something for a (operative) reason, then this acceptance or rejection is performed for good reason only if the (operative) reason is indeed a reason (in our sense) for or against, respectively, what one accepts or rejects. Thus, if one accepts A for the reason Γ, then that Γ implies A is a necessary but not a sufficient condition for one's accepting A being done for good reason.

We have stressed that not all good reasons are logically good reasons, and we call the mistake of thinking otherwise "logicism about reasons." Correspondingly, our implication relation includes more than just logical implications, and our incompatibility relation includes more than just logical incompatibilities. We denote classical propositional consequence by $\vdash_{CL}$. So we know that $\mid\!\sim$ is not identical to $\vdash_{CL}$. Indeed, we have argued that $\mid\!\sim$ ought to differ from classical consequence in fundamental and radical ways. In particular, it should be nonmonotonic and nontransitive. For, our aim is to offer a treatment of logical vocabulary as being able to make explicit open reason relations. Nevertheless, one might plausibly claim that, while not all good implications are logically good implications, all implications of classical logic are indeed good implications, so that $\vdash_{CL} \subset \mid\!\sim$. This classical subset of all implications obeys, of course, traditional structural principles, and this turns out to be the case for our theory below. We are thus not offering a nonmonotonic logic but rather an account of consequence and a logic that can codify generally nonmonotonic material implications as well as monotonic logical implications.

In contrast to standard approaches to nonmonotonic logic, we do not start with a familiar logical consequence relation, like $\vdash_{CL}$, and then try to extend it to a nonmonotonic consequence relation. Rather, we start with a nonmonotonic consequence relation over atomic sentences, which we then extend to a logically complex language in such a way that the result includes classical consequence. The way in which we do this is to take the implications among atomic sentences and to treat them as axioms in a sequent calculus. Moreover, unlike most extant approaches to nonmonotonic consequence, we treat implications and incompatibilities as equally subject to failures of monotonicity.[1] Thus, not only can implications be defeated by adding further premises but also incompatibilities can be "cured" by adding further claims.[2]

The chapter is structured as follows: in the first subsection, we introduce the basic sequent calculus NMMS, which illustrates how logical vocabulary can be elaborated from any autonomous base vocabulary. In the second subsection, we explain two different senses in which the logical vocabulary introduced by NMMS can make reason relations explicit, including open reason relations. Thus, the first two subsections together present a treatment of logical vocabulary as vocabulary that is universally LX, that is, vocabulary that can be elaborated from any autonomous base vocabulary and can make explicit the reason relations among sentences of any base vocabulary from which it is elaborated. In the second subsection, we also formulate a variant of our sequent calculus that works for consequence relations in which Contraction fails. And in the third subsection, we present more adjustments and additions to our sequent calculus. In particular, we show how we can introduce logical vocabulary that makes explicit local structural features of reason relations. The fourth section concludes; and the appendix provides proofs and lemmas that we omit in the main text, in order to make the main text more accessible.

3.1 Sequent Calculi for Logical Expressivism

The brilliant idea behind Gentzen's (1934) invention of sequent calculus is to encode entire implications, rather than just sentences, in objects that are manipulated in a calculus. These objects are called "sequents," and we will write them, for example, like this: $\Gamma \succ \Delta$, where Γ is the set of the premises and Δ is a conclusion or a set of conclusions.[3] In a sequent calculus, instead of deriving sentences from sentences, as in natural deduction or axiomatic systems, we derive sequents from sequents, and the thought is that sequents that are derivable correspond to good implications. This idea fits well with the aims of logical expressivism. For, according to logical expressivism, the proper starting point for the formulation of an account of consequence that includes logical consequence is a set of implications that hold among nonlogical, that is, atomic, sentences. Such implications can serve as the input or axioms of a sequent calculus, as will become clearer below.

Moreover, according to the semantic inferentialism that informs our larger project, such implications and incompatibilities among declarative atomic sentences articulate the conceptual roles played by these sentences. In this sense, the axioms of our sequent calculi articulate the conceptual roles of the nonlogical vocabulary. A sequent calculus extends these conceptual roles of atomic sentences to a consequence relation that also articulates the conceptual roles of logically complex sentences.

3.1.1 How Many Conclusions?

Before moving on, we should pause to note that, for different logics, Gentzen uses two different kinds of sequents, namely sequents that can have at most one conclusion and sequents that can have many conclusions. This is often put by saying that we can work in a Set-Formula framework or in a Set-Set framework, depending on whether the conclusion is a single formula or a set of formulae. And this can be generalized further, in one dimension, by working with lists or multi-sets or even tree-structures instead of sets; and it can be generalized, in another dimension, by working with sequents that don't have two sides but one, or three, or any other number of sides.[4] These differences have important consequences for the logics that are formulated in a sequent calculus. It is, for instance, one of Gentzen's surprising results that his sequent rules yield intuitionistic logic in the Set-Formula framework, but the very same sequent rules yield classical logic in the Set-Set framework.

In the previous chapters, we have often talked as if $\mid\!\sim$ is always flanked by a set of premises on the left and a single conclusion on the right. Thus, we have seemingly used a Set-Formula framework. We did this to avoid confusions that might arise from the fact that it is not immediately clear how to think philosophically about sequents with multiple conclusions. After all, a good implication seems to correspond to a good argument or inference, and such arguments or inferences have one or more premises but just one single conclusion. Similarly, it seems that we give reasons for or against single claims. Note, however, that we said that Γ is a reason (in our technical sense) for A if and only if commitment to accept all the sentences in Γ precludes entitlement to reject A.[5] And Γ is a reason against A if and only if commitment to accept all the sentences in Γ precludes entitlement to accept A. These notions can readily be generalized to allow for multiple conclusions, namely as follows: Γ is a reason for Δ if and only if commitment to accept all the sentences in Γ precludes entitlement to reject all the sentences in Δ. And Γ is a reason against Δ if and only if commitment to accept all the sentences in Γ precludes entitlement to accept all the sentences in Δ.

This way of understanding multiple conclusion sequents is a variant of the view developed by Restall (2005). He calls collections of assertions and denials "positions," and he writes the position in which everything in Γ is asserted and everything in Δ is denied as $[\Gamma : \Delta]$. Restall then suggests that we can understand the claim that Γ implies Δ as the claim that it is normatively improper or "out-of-bounds"—as Restall says—to assert everything in Γ and deny everything in Δ, that is, the position $[\Gamma : \Delta]$ is out-of-bounds. If we think that a position is out-of-bounds just in case one cannot be entitled to all of the commitments in the position, then our

suggestion and Restall's suggestion coincide. We take these to be equivalent ways to understand implication or consequence in pragmatic-normative terms, and we will sometimes refer to this understanding generically as "normative bilateralism."[6]

We will usually work with sequents with multiple conclusions, thus relying on the interpretation of such sequents just given.[7] We usually work with sets of premises and conclusions, but we occasionally also use multisets. So, unless specified otherwise, our sequents will be of the form $\Gamma \succ \Delta$, where Γ and Δ are sets of sentences (availing ourselves of the usual ways to omit set-theoretic notation in sequents). And we think correspondingly about implication or consequence as also allowing for multiple conclusions, so that we can write $\Gamma \mathrel{|\!\sim} \Delta$ to say that the set Γ implies the set Δ or, equivalently, the set Δ is a consequence of the set Γ. As explained above, we will write $\Gamma, \Delta \mathrel{|\!\sim}$ to say that Γ and Δ are incompatible, and so the sentences in one of these sets are jointly a reason against those in the other set (jointly).

3.1.2 Adding Logical Vocabulary

As explained in the previous chapters, we use the term "vocabulary" to mean a collection of expressions together with the reason relations that hold among (sets of) them. And we formulated the core thesis of logical expressivism as the claim that logical vocabulary is, at least as an ideal, vocabulary that is universally LX, that is, it can be *el*aborated from and is *ex*plicative of any autonomous base vocabulary. More precisely, some vocabulary is universally LX just in case, for any base vocabulary that is sufficient to have a discursive practice of giving and asking for reasons for and against claims, (i) the reason relations among sentences that include bits of the universally LX vocabulary can be elaborated from the reason relations that hold among sentences of the base vocabulary, and (ii) the LX vocabulary allows its users to make explicit the reason relations that hold in the base vocabulary (and in the thus extended vocabulary).

In order to formulate a logic in accordance with this idea of logical expressivism, we must consider two questions: (a) What does it mean to elaborate reason relations among sentences featuring logical vocabulary from reason relations among sentences that do not feature any logical vocabulary? (b) What does it mean that sentences featuring logical vocabulary allow us to make explicit reason relations? In the remainder of this section, we address the first of these two questions by formulating an answer in terms of a sequent calculus. We turn to the second question in the following section.

We start with a base or base vocabulary $\mathfrak{B}$, which we assume to be autonomous, in the sense of being sufficient to have a discursive practice

of giving and asking for reasons. We will think of the lexicon of our nonlogical, base vocabulary as a countable set of atomic sentences. Let us call this atomic language $\mathfrak{L}_{\mathfrak{B}}$. We encode the relations of reasons-for and reasons-against among sentences of $\mathfrak{L}_{\mathfrak{B}}$ in a single consequence relation $\mathrel{|\!\sim}_{\mathfrak{B}}$, as explained above. Hence, $\Gamma \mathrel{|\!\sim}_{\mathfrak{B}} \Delta$ holds just in case the sentences in Γ are, jointly, a reason for those in Δ. And $\Gamma, A \mathrel{|\!\sim}_{\mathfrak{B}}$ holds just in case the sentences in Γ are jointly a reason against A.[8]

Definition 1 (Material Base)**.** A material base, $\mathfrak{B}$, is an atomic language, $\mathfrak{L}_{\mathfrak{B}}$, and a base consequence relation, $\mathrel{|\!\sim}_{\mathfrak{B}}$, between subsets of this language $\mathrel{|\!\sim}_{\mathfrak{B}} \subseteq \mathcal{P}(\mathfrak{L}_{\mathfrak{B}}) \times \mathcal{P}(\mathfrak{L}_{\mathfrak{B}})$. A base consequence relation obeys Containment just in case $\langle \Gamma, \Delta \rangle \in \mathrel{|\!\sim}_{\mathfrak{B}}$ whenever $\Gamma \cap \Delta \neq \emptyset$.

In order to extend our base vocabulary so as to include logical vocabulary, we must do two things. First, we must add logical expressions to the lexicon $\mathfrak{L}_{\mathfrak{B}}$. Second, we must specify the reason relations among our thus extended lexicon. We do the first by adding the logical expressions denoted by $\neg$, $\rightarrow$, $\wedge$, and $\vee$ to $\mathfrak{L}_{\mathfrak{B}}$ in the usual way. We call this extended lexicon $\mathfrak{L}$, and, to be precise, we can define it recursively by saying: firstly, all the sentences in $\mathfrak{L}_{\mathfrak{B}}$ are in $\mathfrak{L}$; secondly, if ϕ and ψ are any sentences in $\mathfrak{L}$, then $\neg\phi$, $\phi \rightarrow \psi$, $\phi \wedge \psi$, and $\phi \vee \psi$ are also in $\mathfrak{L}$; and lastly nothing else is in $\mathfrak{L}$.

We turn to a sequent calculus to perform our second task, namely to specify the reason relations among the sentences of our extended lexicon.[9] In formal terms, we have to specify a consequence relation $\mathrel{|\!\sim}$ over sets of sentences of $\mathfrak{L}$, starting from the consequence relation $\mathrel{|\!\sim}_{\mathfrak{B}}$ over $\mathfrak{L}_{\mathfrak{B}}$. Now, while sequent calculi can be formulated in many different ways, one can often formulate them by using sequents with only atomic sentences as axioms. That is what we do here. Where standard approaches use atomic sequents with certain formal properties (such as those with a non-empty intersection of the two sets, that is, instances of Containment), however, we will simply use our entire base consequence relation $\mathrel{|\!\sim}_{\mathfrak{B}}$. So, calling our sequent calculus NMMS (for *n*on-*m*onotonic *m*ulti-*s*uccedent), we stipulate that $\Gamma \succ \Delta$ is an axiom of NMMS if and only if $\Gamma \mathrel{|\!\sim}_{\mathfrak{B}} \Delta$. We write $\mathsf{NMMS}_{\mathfrak{B}}$ if we want to be clear what the base vocabulary is.

We can now use sequent rules to derive sequents with logically complex sentences from our axioms. Sequent rules are usually written as one or more schematic sequents, the top sequents, above a horizontal line and one sequent below that line, the bottom sequent. The sentences that occur unchanged in top sequents and bottom sequents are called the "context," and the sentences that occur, as free-standing sentences, only above or below the line are called the "active formulae" of the rule; sometimes the

active sentence below the line is called the "principal formula." To see how this works, let's look at NMMS:

Axioms of NMMS:
Ax1: $\Gamma \succ \Delta$ is an axiom if and only if $\Gamma \mathrel{\underset{\mathfrak{B}}{\vdash\!\!\sim}} \Delta$.

Rules of NMMS:

$$\frac{\Gamma \succ \Delta, A \qquad B, \Gamma \succ \Delta \qquad B, \Gamma \succ \Delta, A}{\Gamma, A \to B \succ \Delta}\,[\mathrm{L}\to] \qquad \frac{\Gamma, A \succ B, \Delta}{\Gamma \succ A \to B, \Delta}\,[\mathrm{R}\to]$$

$$\frac{\Gamma \succ \Delta, A}{\Gamma, \neg A \succ \Delta}\,[\mathrm{L}\neg] \qquad \frac{\Gamma, A \succ \Delta}{\Gamma \succ \Delta, \neg A}\,[\mathrm{R}\neg]$$

$$\frac{\Gamma, A, B \succ \Delta}{\Gamma, A \wedge B \succ \Delta}\,[\mathrm{L}\wedge] \qquad \frac{\Gamma \succ \Delta, A \qquad \Gamma \succ \Delta, B \qquad \Gamma \succ \Delta, A, B}{\Gamma \succ \Delta, A \wedge B}\,[\mathrm{R}\wedge]$$

$$\frac{\Gamma, A \succ \Delta \qquad \Gamma, B \succ \Delta \qquad \Gamma, A, B \succ \Delta}{\Gamma, A \vee B \succ \Delta}\,[\mathrm{L}\vee] \qquad \frac{\Gamma \succ \Delta, A, B}{\Gamma \succ \Delta, A \vee B}\,[\mathrm{R}\vee]$$

The rules of NMMS are slight variations on rules that are familiar as so-called "Ketonen-style rules" (see Negri et al., 2008; Humberstone, 2011; Bimbó, 2015). The only difference is that the rules that have more than one top sequent have exactly two top sequents in Ketonen-style rules, but they have three top sequents in NMMS. The third top sequent is such that all the active formulae in the other top sequents occur in the third top sequent, and they occur on the same side as in the other top sequent in which they occur. The reason to include the third top sequents in these rules is that they ensure that the number of copies of a sentence on the left or the right of the sequent arrow does not make any difference to which sequents are derivable from which other sequents. The difference matters only when Monotonicity fails. For, Monotonicity allows one to derive the unfamiliar third top sequent from either one of the two traditional top sequents of the Ketonen-style rules.

As an example of what these rules say, the rule [R→] says that if we have a sequent in which a sentence A is on the left side and a sentence B is on the right side, then we can derive the sequent that is like the first except that instead of A on the left and B on the right it includes the sentence $A \to B$ on the right side.[10] Thus, the rule [R→] immediately ensures that one direction of the Deduction-Detachment condition holds. And explanations of the other rules can be given in an analogous way.[11]

A sequent is derivable just in case there is a tree of applications of these sequent rules in which the sequent in question is at the root and every sequent that sits at a leaf of a tree is an axiom. In the following example of such a proof tree, the root is $\succ((p \wedge r) \vee s) \rightarrow q$ and the leaves are $p, r \succ q$ and $s \succ q$ and $p, r, s \succ q$.

$$\dfrac{\dfrac{\dfrac{p, r \succ q}{p \wedge r \succ q}\,[\mathrm{L}\wedge] \qquad s \succ q \qquad \dfrac{p, r, s \succ q}{p \wedge r, s \succ q}\,[\mathrm{L}\wedge]}{(p \wedge r) \vee s \succ q}\,[\mathrm{L}\vee]}{\succ ((p \wedge r) \vee s) \rightarrow q}\,[\mathrm{R}\rightarrow]$$

The sequent calculus NMMS is a way to perform our second task, namely the task of specifying the reason relations among the sentences of our extended lexicon, because we can now say that, in our extended language $\mathfrak{L}$, a set of sentences Γ implies a set of sentences Δ just in case the sequent $\Gamma \succ \Delta$ is derivable in NMMS. That is, $\Gamma \mathrel{|\!\sim} \Delta$ if and only if $\Gamma \succ \Delta$ is derivable in NMMS. In this way, we can take any base vocabulary $\mathfrak{B}$, consisting of a set of sentences $\mathfrak{L}_{\mathfrak{B}}$ and a consequence relation over them $\mathrel{|\!\sim}_{\mathfrak{B}}$, and extend it to include logical vocabulary, by moving to the language $\mathfrak{L}$ and consequence relation $\mathrel{|\!\sim}$ over the extended language $\mathfrak{L}$. The upshot of the example proof tree above, for example, is that if $p, r \mathrel{|\!\sim}_{\mathfrak{B}} q$ and $s \mathrel{|\!\sim}_{\mathfrak{B}} q$ and $p, r, s \mathrel{|\!\sim}_{\mathfrak{B}} q$, then $\mathrel{|\!\sim} ((p \wedge r) \vee s) \rightarrow q$. We call the reason relations that hold among sentences in our logically extended lexicon the "logically extended reason relations."

Note that NMMS does not include any of the following rules which are often used in sequent calculi and are among the traditional structural principles (together with Permutation, which we take for granted throughout this book), where the first rule allows us to derive a bottom sequent without using any top sequents:

$$\dfrac{}{A \succ A}\,[\text{Reflexivity}]$$

$$\dfrac{\Gamma, A, A \succ \Delta}{\Gamma, A \succ \Delta}\,[\text{L-Contraction}] \qquad \dfrac{\Gamma \succ A, A, \Delta}{\Gamma \succ A, \Delta}\,[\text{R-Contraction}]$$

$$\dfrac{\Gamma \succ \Delta}{\Gamma, \Theta \succ \Delta, \Lambda}\,[\text{Weakening}] \qquad \dfrac{\Gamma \succ \Delta, A \qquad A, \Theta \succ \Lambda}{\Gamma, \Theta \succ \Delta, \Lambda}\,[\text{Mixed-Cut}]$$

The reason why NMMS does not include [Reflexivity], [L-Contraction], and [R-Contraction] as rules is very different from the reason why NMMS does not include [Weakening] and [Mixed-Cut]. It doesn't include the former rules because, given a base that obeys Containment, these rules are not necessary; adding them to NMMS does not change what is derivable in NMMS. All the sequents derivable by [Reflexivity] are already axioms if the

sentence in the sequent is atomic, and one can derive them using the rules of NMMS if the sentence is logically complex. The sequent $p \wedge q \succ p \wedge q$, for instance, is derivable from $p, q \succ p$ and $p, q \succ q$ and $p, q \succ p, q$, by using [R∧] and [L∧]; and it is easy to show that this is always possible (see Appendix, Proposition 24). The two contraction rules are not rules of NMMS because we are working with sets of sentences on the left and the right side of sequents. Hence, the top sequents and the bottom sequents of the contraction rules are identical (although this will not be true for the calculus $\mathsf{NMMS}^{\backslash ctr}$ below). Thus, the rules for Reflexivity and Contraction are not included because they would be redundant and we prefer not to list them as rules.

The reason why NMMS does not include [Weakening] and [Mixed-Cut] is very different. The [Weakening] rule encodes the constraint that reason relations must be monotonic, and the [Mixed-Cut] rule is a generalized version of the idea that reason relations must be transitive. These rules fail and have counterexamples in NMMS when we codify open reason relations, assuming that our sequent rules ought to extend the base consequence relation in a conservative way, that is, assuming that if all the sentences in Γ and Δ are in the base lexicon, then $\Gamma \vdash\!\sim \Delta$ just in case $\Gamma \mathrel{\vdash\!\sim_{\mathfrak{B}}} \Delta$ (see below). For [Weakening], a base consequence relation that is such that $p \mathrel{\vdash\!\sim_{\mathfrak{B}}} q$ but $p, r \mathrel{\not\vdash\!\sim_{\mathfrak{B}}} q$ is a counterexample. And [Mixed-Cut] actually implies [Weakening], given Containment, and so it yields to the same counterexample. To see this, suppose again that $p \mathrel{\vdash\!\sim_{\mathfrak{B}}} q$. Now, by Containment, $p, q, r, \mathrel{\vdash\!\sim_{\mathfrak{B}}} q$. And so [Mixed-Cut] would allow us to prove that $p, r \vdash\!\sim q$. Hence, we must reject [Weakening] and [Mixed-Cut] and allow that they can fail. These failures make the consequence relations that we shall study substructural: some of the traditional structural principles fail in them.

Returning to the thesis of logical expressivism, note that vocabulary that can be introduced into a language in the way in which we have just introduced the vocabulary of propositional logic can be elaborated from any autonomous base vocabulary, and it thus meets the first part of the condition of logicality suggested by logical expressivism, that is, the first part of the condition that logical vocabulary must be universally LX. For, someone who can use any base language, in accordance with its reason relations, already knows how to do everything they need to know how to do in order to acquire the ability to use the vocabulary introduced in this subsection, in accordance with its reason relations. All that such a person needs to do is to leverage their sensitivity to reason relations among atomic sentences to yield a sensitivity to reason relations among sentences that include logical vocabulary in accordance with the sequent rules above.[12] And we can leave it open what exactly the relevant kind of sensitivity to

reason relations is that a speaker must have in order to be a competent speaker of some base language. As long as this ability can be transformed in a way that corresponds to the sequent rules above, all that a speaker needs to do in order to acquire the ability to use logical vocabulary is to undergo such a transformation. Crucially, this does not require any new experience of, or familiarity with, objects or properties that are novel to the speaker. In this sense, the potentiality of understanding and using logical vocabulary is contained in the ability to understand and use any autonomous vocabulary whatsoever. This is the first part of the idea that logical vocabulary is vocabulary that is universally LX, and we have now explained how extending a base vocabulary by means of a sequent calculus in the way sketched above meets this first criterion of logicality.

3.1.3 Some Familiar Features of Logical Vocabulary

The result of adding logical vocabulary and specifying its reason relations by NMMS is a radical departure from standard approaches to logic in many ways, but it is also conventional in other ways. So let us highlight some properties of reason relations that are extended by NMMS. We start, in this subsection, with features that our theory shares with many familiar sequent calculi treatments of logical vocabulary, and we will discuss the radical and novel aspects of our theory in the next subsection.

As already intimated, the rules of NMMS are slight variations on rules that are familiar as so-called "Ketonen-style rules" (see Negri et al., 2008; Humberstone, 2011; Bimbó, 2015). And the rules of NMMS share the key virtues of the Ketonen-style rules, in particular the following three: First, the rules allow us to derive all classically valid sequents without the help of structural rules if the base consequence relation obeys Containment. Second, the rules are such that every sentence that occurs in any of the top sequents of a rule application also occurs, either as a free-standing sentence or as a subformula of a sentence, in the bottom sequent,[13] and the top sequents are always all at most as logically complex as the bottom sequent. Third, the rules are invertible, which means that if the bottom sequent of an application of a rule is derivable, then all the top sequents are also derivable. Let us explain these three features in a bit more detail.

Regarding the first familiar feature of NMMS, note that if the base consequence relation obeys Containment, then any sequent of the form $\Gamma, p \succ p, \Delta$, of sentences in the base lexicon, is an axiom. And it can easily be shown that closing sequents of this form under the rules of NMMS yields exactly the relation of classical propositional consequence, over the extended lexicon. Since adding more sequents as axioms does not affect the derivations of the classically valid sequents, closing any base consequence

relation that obeys Containment under the rules of NMMS includes all the implications of classical propositional logic.

Fact 2. *If a base obeys Containment, then any implication among sentences of $\mathfrak{L}$ that holds in classical propositional logic also holds in $\mid\!\sim$, that is, $\vdash_{CL} \subseteq \mid\!\sim$. (Appendix, Proposition 25)*

In other words, insofar as everything implies itself (with arbitrary contexts), we agree with classical logicians regarding every classically valid implication that it is indeed valid. And since almost all non-classical logics are strictly weaker than classical logic,[14] in that they deny the validity of some implication that is deemed valid by classical logic, our position can seem very classical at this point. Indeed, if we say that the "narrowly logical part" of the logically extended consequence relation is the part that can be derived in NMMS from just the instances of Containment, then the narrowly logical part of our theory just is classical logic. Equivalently, the logically good reason relations are those that hold in the logical extensions of all base vocabularies that obey Containment. Of course, all of the structural principles hold in the narrowly logical part of our logically extended consequence relation; in particular, this narrowly logical part obeys Monotonicity and Transitivity. So, while the logically extended consequence relation of a base that obeys Containment will typically be nonmonotonic and nontransitive, the narrowly logical part of this consequence relation will be monotonic and transitive. In this sense, we are offering a fully structural logic for the purpose of codifying open reason relations, that is, reason relations that are substructural. Thus, if we identify a logic with the narrowly logical part of a consequence relation, then we are not offering a nonmonotonic logic, but rather a monotonic—indeed classical—logic whose logical vocabulary allows us to make explicit nonmonotonic consequence relations.

The consequence relations defined by NMMS are substructural in the sense that Monotonicity and Transitivity can fail in them, namely if they fail in the base consequence relations that are extended by the rules of NMMS. Sequent rules for the connectives that are equivalent in the context of structural rules often differ in important respects in substructural settings. The so-called multiplicative and additive rules of linear logic are perhaps the most common rules for conjunction and disjunction in a substructural setting, thus yielding two conjunctions and two disjunctions. Our (quasi) Ketonen-style sequent rules use, in effect, additive rules on one side and multiplicative rules on the other side. If our base consequence relations obey Containment, then the differences between these rules only matter outside of the narrowly logical part of the consequence relation. In addition to the features of the Ketonen rules that we point out below, one important

reason to adopt the Ketonen rules is that they allow us to respect two plausible constraints, which we will illustrate only for conjunction (but which have analogues for disjunction). The first constraint is that if the addition of the premise "Tweety is a penguin" defeats the implication from "Tweety is a bird" to "Tweety can fly," then replacing the premise "Tweety is a bird" with the conjunctive premise "Tweety is a bird, and Tweety is a penguin" should not imply "Tweety can fly." The additive conjunction left-rule would not allow us to respect this first constraint. The second constraint is that it can happen, for instance, that "Tweety is a bird" implies "Tweety can fly" and "Tweety is a plastic toy" implies "Tweety is inanimate," but the combination of the premises "Tweety is a bird" and "Tweety is a plastic toy" does not imply "Tweety can fly, and Tweety is inanimate." However, the multiplicative conjunction right-rule would not allow us to respect this constraint. By using the (quasi) Ketonen-style rules, we can respect both constraints. Thus, our rules allow us to remain classical within the narrowly logical part of the consequence relation, while also allowing us to respect the two just stated constraints outside of the narrowly logical part of the consequence relation.

The second familiar feature of the rules of NMMS is that in moving towards the root on any branch of a proof tree, the sequents never get logically less complex[15] and never fail to include a sentence (embedded or otherwise) that already occurred in any sequent in the branch. Thus, in building a proof tree we compose the root sequent out of all and only the material (that is, atomic sentences) provided by the leaves of the tree. This feature is not only very useful for proving results about sequent calculi, it also lets us see how logically complex sequents derive from our base consequence relation. Moreover, it is an immediate consequence of this feature that the logical extension of any base vocabulary by NMMS is conservative: if all the sentences in Γ and Δ are in our base lexicon $\mathfrak{L}_{\mathfrak{B}}$, then $\Gamma \mid\!\sim \Delta$ just in case $\Gamma \mid\!\sim_{\mathfrak{B}} \Delta$.

Fact 3. *The extension of any consequence relation* $\mid\!\sim_{\mathfrak{B}}$ *to* $\mid\!\sim$ *by* NMMS *is a conservative extension: if* $\Gamma \cup \Delta \subseteq \mathfrak{L}_{\mathfrak{B}}$, *then* $\Gamma \mid\!\sim \Delta$ *just in case* $\Gamma \mid\!\sim_{\mathfrak{B}} \Delta$. *(Appendix, Proposition 26)*

This is an important fact from the perspective of logical expressivism because if logical vocabulary has the function of making explicit reason relations in a base vocabulary, then its introduction into the language should not add or subtract any reason relations that hold between sentences of the base language. Logical vocabulary should make explicit and not change the reason relations with which we started. That our way of adding

logical vocabulary is a conservative extension of our base consequence relation ensures this.

The third familiar feature of NMMS is that the rules are invertible (see Appendix, Proposition 27). That means that we can move from the bottom sequent of a rule application to any of the top sequents. If, for instance, we know that $A \vee B \succ C$ is derivable, then we may conclude that $A \succ C$ is derivable, by applying the inverted rule [L$\vee$]. This use of the rule is not itself a derivation, but it is nevertheless the case that the bottom sequent of an application of [L$\vee$] is derivable only if each top sequent is derivable. This is usually expressed by saying that the inverted rule is *admissible* but not *derivable*. All the rules of NMMS are invertible in this sense. The invertibility of our rules has many consequences, such as the following, of which we already know the first two as conditions of adequacy for expressivist logical theories from the previous chapter:

Deduction-Detachment (DD) Condition on Conditionals:
$\Gamma \mid\!\sim A \to B$ if and only if $\Gamma, A \mid\!\sim B$.

Incoherence-Incompatibility (II) Condition on Negation:
$\Gamma \mid\!\sim \neg A$ if and only if $\Gamma, A \mid\!\sim$ if and only if $\Gamma \# A$.

Antecedent-Adjunction (AA) Condition on Conjunctions:
$\Gamma, A \wedge B \mid\!\sim \Delta$ if and only if $\Gamma, A, B \mid\!\sim \Delta$.

Succedent-Summation (SS) Condition on Disjunctions:
$\Gamma \mid\!\sim A \vee B, \Delta$ if and only if $\Gamma \mid\!\sim A, B, \Delta$.

Thanks to the invertibility of the rules of NMMS, our theory underwrites all of these conditions. As we will see in the next section, that these conditions all hold in our theory is one sense in which the logical vocabulary that is introduced by the rules of NMMS allows us to make explicit reason relations.

To sum up, given any base consequence relation that obeys Containment, its logical extension includes all classically valid implications, it is a conservative extension of the base consequence relation, and it obeys the Deduction-Detachment, the Incoherence-Incompatibility, the Antecedent-Adjunction, and the Succedent-Summation Conditions.

3.1.4 Radical Novelties

In the previous subsection, we have seen that our logical extensions of base consequence relations yield results that are familiar in many ways. Let us now turn to the ways in which our theory is new and unusual—indeed, radically so.

Our account is designed to capture reason relations that are open in the sense that they do not obey structural rules like Monotonicity and Transitivity, that is, the rules of [Weakening] and [Mixed-Cut]. A base consequence relation may, for instance, plausibly include and not include, respectively, the following implications, which are jointly an example of failure of Monotonicity (in which we label premises and conclusions in square brackets as we go along, for future reference)[16]:

(1) [*p*] Tara is a human being. $\mathrel{|\!\sim}_{\mathfrak{B}}$ [*q*] Tara's body temperature is 37 degrees Celsius.

(2) [*p*] Tara is a human being. [*r*] Tara has a fever. $\not\mathrel{|\!\sim}_{\mathfrak{B}}$ [*q*] Tara's body temperature is 37 degrees Celsius.

Here (1) is a good implication, but it is defeated by the additional premise that Tara has a fever. We may call such a premise a "defeating premise" or a "defeater." Of course, one might object that (1) is not a good implication because the truth of the premise does not guarantee the truth of the conclusion, as is brought out by (2). Such an objection is, however, merely a way to express the unwillingness to consider the possibility of open reason relations. Once we reject logicism about reasons, it is hard to see what could justify such an unwillingness. So we set this worry aside.

Note that monotonicity may fail not only for implications, or reasons-for, but also for incompatibilities, or reasons-against. Here is an example that we might want to include in our base consequence relation.

(3) [*s*] This figure is a triangle. [*t*] The sum of the inner angles of this figure is larger than two right angles. $\mathrel{|\!\sim}_{\mathfrak{B}}$

(4) [*s*] This figure is a triangle. [*t*] The sum of the inner angles of this figure is larger than two right angles. [*u*] This figure is a spherical triangle. $\not\mathrel{|\!\sim}_{\mathfrak{B}}$

Here we have two sentences that are incompatible by themselves, but they become compatible in the presence of the additional sentence that the figure in question is a spherical triangle. We may also express this phenomenon by saying that the incoherence of the set of the first two sentences is cured by the addition of the third sentence.

Moreover, such examples of nonmonotonicity can be turned into examples in which transitivity fails.[17] In particular, if we look at a sentence that is close in its meaning to the negation of the defeating premise and also close to the defeated conclusion, we can construct counterexamples to transitivity by considering these sentences as conclusions. For our purposes the sentences "Tara does not have a fever" and "Tara is healthy" are close enough to allow us to construct the following example:

(5) [p] Tara is a human being. $\underset{\mathfrak{B}}{\vdash\!\!\sim}$ [q] Tara's body temperature is 37 degrees Celsius.

(6) [p] Tara is a human being. [q] Tara's body temperature is 37 degrees Celsius. $\underset{\mathfrak{B}}{\vdash\!\!\sim}$ [v] Tara is healthy.

(7) [p] Tara is a human being. $\underset{\mathfrak{B}}{\not\vdash\!\!\sim}$ [v] Tara is healthy.

If we were allowed to apply [Mixed-Cut], we could derive (7) from (5) and (6). But while (5) and (6) are intuitively good implications, it seems implausible that if we are committed to accept that Tara is a human being, we cannot be entitled to reject that Tara is healthy, that is, (7) is intuitively not a good implication. Hence, we have an intuitive counterexample to [Mixed-Cut].

We can also apply this strategy to failures of monotonicity among incompatibilities, if we have sentences that behave closely enough to negations of the sentences in our example of monotonicity failures. Here is how we can apply this to the example (3)-(4) above:

(8) [s] This figure is a triangle. $\underset{\mathfrak{B}}{\vdash\!\!\sim}$ [w] The sum of the inner angles of this figure is equal to two right angles.

(9) [s] This figure is a triangle. [w] The sum of the inner angles of this figure is equal to two right angles. $\underset{\mathfrak{B}}{\vdash\!\!\sim}$ [x] This figure is a Euclidean plane triangle.

(10) [s] This figure is a triangle. $\underset{\mathfrak{B}}{\not\vdash\!\!\sim}$ [x] This figure is a Euclidean plane triangle.

Perhaps an opponent would want to say that (10) is indeed a good implication because a triangle is by definition a plane figure and a spherical triangle is, hence, not a triangle. Or an opponent might object that as long as we count spherical triangles as triangles, implication (8) is not a good implication. Of course, it is useful in mathematics to have exact definitions that work in ways that underwrite such objections. But that just means that it is useful in mathematics to have exact definitions that allow our reason relations to be monotonic. We wholeheartedly agree with that. Nothing here hangs on any particular example, and the way we use "triangle" in our example is the way in which spherical triangles are neither excluded nor prominent enough in the concept expressed by "triangle" to undermine the defeasible implication (8), which we take to be the everyday use of the term that is regimented in exact ways in mathematics.[18]

The aspect of our theory that is probably most radical and new is that we take examples like those just given at face value and include them in

our consequence relations. We do not try to reconstruct them, on the basis of a monotonic logic, with the help of default rules or partial orderings of models, like prominent approaches to nonmonotonic logics do. Rather, we take the collection of all implications among atomic sentences, those of the examples just given included, as our base consequence relation and we, thus, treat them as axioms of NMMS.

As a toy example, we can do this for the implications above. To do so, we treat the sentences p,...,x as the only atomic sentences of our base lexicon, and we say that our base consequence relation are exactly the implications described as good above plus all instances of Containment. We may call this base $\mathfrak{T}$ (for "toy") and specify it thus:

$$\mathfrak{L}_{\mathfrak{T}} = \{p, q, r, s, t, u, v, w, x\}$$
$$\mathrel{|\!\sim}_{\mathfrak{T}} = \{\langle\{p\},\{q\}\rangle, \langle\{s,t\},\emptyset\rangle, \langle\{p,q\},\{v\}\rangle, \langle\{s\},\{w\}\rangle, \langle\{s,w\},\{x\}\rangle\} \cup \{\langle\Gamma,\Delta\rangle : \Gamma \cup \Delta \subseteq \mathfrak{L}_{\mathfrak{T}} \text{ and } \Gamma \cap \Delta \neq \emptyset\}$$

The part of the definition of $\mathrel{|\!\sim}_{\mathfrak{T}}$ that comes after the first union sign ensures that $\mathrel{|\!\sim}_{\mathfrak{T}}$ obeys Containment. So, it follows from our observations above that the logical extension of this base includes all classically valid inferences (over the language that results from adding the logical lexicon to $\mathfrak{L}_{\mathfrak{T}}$). In addition to these classically valid implications, however, $\mathrel{|\!\sim}_{\mathfrak{T}}$ and its logical extension $\mathrel{|\!\sim}$ include the implications that correspond to (1), (3), (6), (8), and (9) above. And because the extension is conservative, neither consequence relation includes the implications that correspond to (2), (4), (7), or (10).

In addition to the implications of our base consequence relation, the logically extended consequence relation includes implications of logically complex sentences that make explicit the good implications of (1), (3), (6), (8), and (9) above, which we may label with the subscript e for "explicitation": (1_e) $\mathrel{|\!\sim} p \rightarrow q$, ($3_e$) $\mathrel{|\!\sim} \neg(s \wedge t)$, ($6_e$) $\mathrel{|\!\sim} (p \wedge q) \rightarrow v$, ($8_e$) $\mathrel{|\!\sim} s \rightarrow w$, ($9_e$) $\mathrel{|\!\sim} (s \wedge w) \rightarrow x$. The sentences on the right in these implications are all theorems of the logically extended consequence relation, in the sense of being implications of the empty set. None of them is, however, a theorem of classical logic. Rather, that these sentences are theorems of the extended consequence relation makes explicit what implies what and what is incompatible with what in our base consequence relation. For example, (1_e) makes explicit, in the form of a theorem, that p implies q. And (3_e) makes explicit, in the form of a theorem, that s and t are incompatible.

Let us highlight, as a side remark, that what matters here is not the status of these sentences as theorems. As will become clearer below, these sentences could play their explicitating role equally well in the context

of other premises and conclusions. And more generally, the notion of a theorem, as a sentence that follows from the empty set, loses much of its interest once we allow that the sentence may not follow from other premises. The more interesting notion is that of a sentence that follows from all premise-sets. Given a base that obeys Containment, this is true in NMMS of all the theorems of classical logic; but a logical extension can also include other sentences of which this is true. We may call such sentences "persistent theorems." We will return to the issue of persistence in the third section of this chapter. At present, we can ignore it and simply note that the theorems in the logically extended vocabulary that are listed above reflect material reason relations that hold in our base consequence relation.

Many familiar principles fail outside of the narrowly logical part of the logically extended consequence relation. The examples from above show that [Weakening] (MO) and [Mixed-Cut] (CT) fail:

MO-failure-1: $p \mathrel{|\!\sim} q$. But $p, r \not\mathrel{|\!\sim} q$.

MO-failure-2: $s, t \mathrel{|\!\sim}$. But $s, t, u \not\mathrel{|\!\sim}$.

CT-failure-1: $p \mathrel{|\!\sim} q$. And $p, q \mathrel{|\!\sim} v$. But $p \not\mathrel{|\!\sim} v$.

CT-failure-2: $s \mathrel{|\!\sim} w$. And $s, w \mathrel{|\!\sim} x$. But $s \not\mathrel{|\!\sim} x$.

A moment's reflection shows that what is sometimes called "meta-*modus-ponens*" can also fail in the logical extensions of base consequence relations. Meta-*modus-ponens* says that if ϕ is a theorem and $\phi \rightarrow \psi$ is a theorem, then ψ is a theorem. Now, we can change our examples of failures of CT by replacing the premise of the first implication by the empty set, which yields sequents of the form $\mathrel{|\!\sim} \phi$, and $\phi \mathrel{|\!\sim} \psi$. But $\not\mathrel{|\!\sim} \psi$. Since $\phi \mathrel{|\!\sim} \psi$ if and only if $\mathrel{|\!\sim} \phi \rightarrow \psi$, cases of this schematic form are cases in which meta-*modus-ponens* fails. Note, moreover, that such failures can occur despite the fact that, as a classically valid schema, all instances of *modus ponens* hold, that is, all instances of $\phi, \phi \rightarrow \psi \mathrel{|\!\sim} \psi$ hold and, indeed, they are all indefeasible, that is, all results of applications of [Weakening] to such instances also hold. In other words, if the premises of a *modus ponens* inference are merely implied, then the conclusion of the *modus ponens* does not follow; but if the premises are explicitly contained in the premises, then the conclusion follows.

In order to highlight how radically our theory diverges from traditional theories outside of the narrowly logical part of the extended consequence relations, let us look at one last example that combines failures of CT, meta-*modus-ponens*, and distribution principles as substitution rules. Suppose we have a base consequence relation that obeys Containment and that includes the following implications: $\mathrel{|\!\sim}_{\mathfrak{B}} p$; $\mathrel{|\!\sim}_{\mathfrak{B}} q, r$; $\mathrel{|\!\sim}_{\mathfrak{B}} p, q, r$. Let us also

stipulate, however, that it is not the case that $\mathrel{|\!\sim}_{\mathfrak{B}} p, r$. These implications allow us to construct the following proof tree.

$$\dfrac{\succ p \qquad \dfrac{\succ q, r}{\succ q \vee r}\,[\mathrm{R}\vee] \qquad \dfrac{\succ p, q, r}{\succ p, q \vee r}\,[\mathrm{R}\vee]}{\succ p \wedge (q \vee r)}\,[\mathrm{R}\wedge]$$

Hence, we have $\mathrel{|\!\sim} p \wedge (q \vee r)$. Furthermore, the following is a classical valid implication and hence part of our logically extended consequence relation: $p \wedge (q \vee r) \mathrel{|\!\sim} (p \wedge q) \vee (p \wedge r)$ and $\mathrel{|\!\sim} (p \wedge (q \vee r)) \rightarrow ((p \wedge q) \vee (p \wedge r))$. However, our extended consequence relation does not include the sequent $\mathrel{|\!\sim} (p \wedge q) \vee (p \wedge r)$. To see this, consider what a proof tree of this sequent would look like. We would have to derive $\succ (p \wedge q), (p \wedge r)$, which requires in turn that we derive $\succ p, (p \wedge r)$. And a derivation of that sequent would have to look as follows:

$$\dfrac{\succ p, p \qquad \succ p, r \qquad \succ p, p, r}{\succ p, p \wedge r}\,[\mathrm{R}\wedge]$$

Since the middle top sequent of this rule application contains only atomic sentences, it holds just in case $\mathrel{|\!\sim}_{\mathfrak{B}} p, r$. But we know that $\not\mathrel{|\!\sim}_{\mathfrak{B}} p, r$. Hence, $\not\mathrel{|\!\sim} (p \wedge q) \vee (p \wedge r)$. This example illustrates, firstly, that CT and meta-*modus-ponens* can fail in our logically extended consequence relation, as we have already seen above. Secondly, the example illustrates that replacing classically equivalent sentences for one another in an implication can turn a good implication into a bad implication. For $p \wedge (q \vee r)$ and $(p \wedge q) \vee (p \wedge r)$ are classically equivalent sentences, but while $\mathrel{|\!\sim} p \wedge (q \vee r)$ holds in our example $\mathrel{|\!\sim} (p \wedge q) \vee (p \wedge r)$ does not hold. And to fully appreciate this point, note that the sentences are not just classically equivalent but they also imply each other in the extended consequence relation in which their substitution as conclusions fails.[19]

To sum up, we have seen in the previous subsection that the logical extensions of base consequence relations that result from applying the rules of NMMS yield results that are familiar and unsurprising in many ways, for instance, by including all classically valid implications. In this subsection, we have seen how radically these logically extended consequence relations diverge from standard treatments of logical vocabulary. These consequence relations can encode open reason relations and, therefore, they are substructural, allowing for failures of Monotonicity and Transitivity. They include implications and theorems that are material in the sense of not being logically valid or logically true, respectively. Principles like meta-*modus-ponens* can fail in such extended consequence relations, even though *modus ponens* holds indefeasibly. And substituting sentences that imply each other

in classical logic (and so also in our logical extensions) for one another in an implication can turn a good implication into a bad one, that is, such substitutions are not guaranteed to preserve consequence. All of these unfamiliar and radical features of our logically extended consequence relations are a result of including nonlogical, material implications in these consequence relations, in particular substructural material implications. In other words, all of these features are consequences of the ability to codify material and open reason relations.

3.2 Making Reason Relations Explicit

In the previous section, we have seen how we can extend reason relations among sentences of an atomic base language to reason relations among the logical extension of that language. Thus, we have seen how logical vocabulary can be elaborated from arbitrary material reason relations. This was the first part of the central claim of logical expressivism, namely the claim that logical vocabulary is universally LX. The second part—the X-part—of this claim is that logical vocabulary allows its users to make explicit the reason relations that hold in the base vocabulary as well as those that hold in the logically extended vocabulary. In this section, we want to clarify what it means for a vocabulary to allow us to make reason relations explicit. Doing this will also give us occasion to introduce some extensions and variants of the sequent calculus NMMS.

3.2.1 The Idea of Making Reason Relations Explicit

The basic idea of making reason relations explicit is to turn something that one acknowledges in practice into something that one acknowledges explicitly in what one says. It is a genus of the species of explicitation that takes as its starting point a practical ability, such as cooking or playing a musical instrument, and yields as its result sentences that convey the norms or rules that govern the exercise of the abilities, such as a recipe for a dish or instructions for playing an instrument. If one can cook well or play a musical instrument well, one has the ability to acknowledge the norms or rules of cooking or playing the instrument in one's practice by non-accidentally acting in accordance with them. When one formulates this practical knowledge of the norms or rules as sentences of a recipe or instructions, one acknowledges them explicitly. One makes these norms or rules explicit. Someone can have the ability to acknowledge such norms or rules in practice without having the ability to make them explicit in the form of principles. And someone might be able to convey such norms or rules explicitly while being much less able to acknowledge them in the practical sense of acting non-accidentally in accordance with them.

The ability to make explicit what one acknowledges in one's practice is crucial for engaging in critical reflection and discussion of what the correct norms or rules governing the activity in question are. If we want to engage in critical reflection and discussion about what it is to cook well or to play an instrument well, for instance, we must be able to make explicit what someone who cooks well or plays an instrument well acknowledges practically in what they do.

According to logical expressivism, there is a particular species of the genus of explicitation for which the ability at issue is the ability to assess and respond to reasons for and against claims, which we may call the ability to practically appreciate reason relations. This species of explicitation takes as its starting point the ability to practically appreciate reason relations and yields as its result sentences that convey the norms or rules that govern the practice of practically appreciating reason relations. It makes explicit reason relations. It gives us sentences that allow us to convey what is a reason for or against what. And just as making the abilities to cook or to play an instrument explicit allows us to engage in critical reflection and discussion about exercises of the abilities to cook or play an instrument, so making explicit reason relations allows us to engage in critical reflection and discussion about exercises of the ability to practically appreciate reason relations. According to logical expressivism, it is the essential function of logical vocabulary to provide us with sentences that allow us to do this, that is, sentences that make explicit what is a reason for or against what. Using logical vocabulary, reason relations of implication and incompatibility can be made explicit in claims, and so can themselves be challenged and defended.

One special feature of the species of explicitation that uses logical vocabulary is that the activity of practically appreciating reason relations and the activity of engaging in critical reflection and discussion regarding reason relations are the same kind of activity. For engaging in critical reflection and discussion is one way to engage in practically appreciating reason relations. Hence, we are here dealing with the possibility that exercises of our conceptual abilities make explicit the norms or rules that govern exercises of these very conceptual abilities. In this sense, the explicitation that is at issue in logic is a case in which our conceptual abilities are directed at themselves; it is a kind of self-consciousness.

A related second special feature of this species of explicitation is that it happens in the same language in which the explicitated practical appreciation of reason relations takes place. When we make reason relations explicit by using logical vocabulary, we do not use an overt meta-language in which we make explicit reason relations among sentences of an object-language. When we assert, for instance, the sentence $A \rightarrow B$ and treat it as undeniable (relative to a given context), we are not saying

that A is a reason for B (relative to the context). Rather, we say that $A \rightarrow B$; but what we do in saying this is to commit ourselves to A being a reason for B (relative to the context). Thus, logical expressivism must not be misunderstood as primarily being a claim about what logically complex sentences say. Rather, it is primarily a claim about what we can do with logically complex sentences, namely to make explicit reason relations. This essential role of logically complex sentences constrains their inferential roles. And according to semantic inferentialism, we can understand the content of logically complex sentences in terms of their inferential roles. Nevertheless, according to logical expressivism, logically complex sentences do not mention but use their constituent sentences; and to assert a logically complex sentence is not to describe some reason relation as holding, but to endorse the reason relation that the sentence makes explicit.

Moreover, according to logical expressivism, logic has a special status among all intellectual pursuits in general and in philosophy in particular. Logic is universal in its scope in that (ideally) it allows us to make explicit any reason relations whatsoever. And these reason relations are what we appreciate in practice in any intellectual pursuit. To be self-conscious about any intellectual pursuit is to be able to engage in critical reflection and discussion about it. So, the ability to engage in logic and use logical vocabulary is an aspect of our self-consciousness; it is the ability to engage in critical reflection and discussion regarding any intellectual pursuit whatsoever. Philosophy is not only crucially concerned with what is a reason for or against what in many particular domains, but critical reflection and discussion of the reason relations that are practically appreciated in doing philosophy are essential to philosophy itself. In this sense, philosophy is an essentially self-conscious intellectual pursuit. And logic studies the vocabulary by which we can manifest this kind of self-consciousness. Thus, logic has a special position with respect to philosophy in particular, because it studies and enables us critically to control the use of the vocabulary that allows us to engage in philosophy self-consciously, that is, to do philosophy with the ability to engage in critical reflection and discussion regarding philosophy. And this ability is an essential part of philosophy, and indeed of any sophisticated intellectual pursuit at all.

We can now see that our question about what it means to make reason relations explicit is really asking what we need in order to engage in critical reflection and discussion regarding reason relations, and so the meanings of our nonlogical words. We argued in the previous chapter that at a minimum this requires that we have a conditional that obeys DD and a negation that obeys II. And we saw in the previous section that the logical vocabulary introduced by NMMS meets these desiderata. We can now spell out the idea of explicitation that lies behind DD and II. It turns out that this can be

done in two different ways. We discuss the first of these ways in the next subsection and then turn to the second way to spell out the idea.

3.2.2 Explicitation by Implication

Let us look at a schematic example of a case where the explicitating function of logical vocabulary becomes relevant. Suppose my, of course fallible, ability to practically appreciate reason relations includes a disposition to practically treat A as a reason for B and to treat A together with B as a reason against C. Thus, we may say that in my exercises of the ability to appreciate reason relations, I act as if $A \mid\!\sim B$ and as if $\{A, B\}\#C$, that is, as if $A, B, C \mid\!\sim$. Suppose that neither of us is willing to accept or reject any of A, B, or C. However, you are not willing to treat A, supposing we accepted it, as a reason for B, nor are you willing to treat A together with B as a reason against C. For you, these three sentences have nothing to do with one another, as far as reason relations go. Finally, suppose that we have (and take ourselves to have) shared background commitments, namely we accept everything in Γ and we reject everything in Δ. What would we need in order to start to discuss and reflect on our disagreement?

Note that we cannot engage in any critical discussion or reflection by wondering about A or B or C directly. For we want to engage in critical reflection about what we would have reason to accept or reject before and independently of engaging in any particular ground-level commitments. We would have some way forward, however limited it may be, if there were a sentence such that treating this sentence as undeniable covaried with treating A as a reason for B and if there was another sentence such that treating that sentence as undeniable covaried with treating A together with B as a reason against C. It is easy to see that, according to NMMS, $\Gamma, A \mid\!\sim B, \Delta$ if and only if $\Gamma \mid\!\sim A \rightarrow B, \Delta$; and we have $\Gamma, A, B, C \mid\!\sim \Delta$ if and only if $\Gamma \mid\!\sim \neg(A \wedge B \wedge C), \Delta$. So, the logically complex sentences $A \rightarrow B$ and $\neg(A \wedge B \wedge C)$ fit the bill, respectively. Given my dispositions to acknowledge reason relations in practice, I have to treat these sentences as undeniable, relative to our shared background commitments, whereas you will refuse to treat them as undeniable. If we have the ability to dispute and reflect on whether a sentence is undeniable, we now have sentences that can serve as the targets in our disagreement. In this sense, the sentence $A \rightarrow B$ allows us to make explicit that A is a reason for B, and the sentence $\neg(A \wedge B \wedge C)$ allows us to make explicit that C is incompatible with A together with B—and thereby bring these reason relations within reach of our critical practices of challenging and defending claims.

In order to serve as the target of our dispute, what matters is not whether you or I accept $A \rightarrow B$. Rather, what matters is whether we treat rejecting it as ruled-out, as out-of-bounds, that is, whether we treat it as if one cannot

be entitled to such a rejection. In other words, what matters is whether we treat our background commitments as providing reasons for $A \to B$. I will do so, and you won't. We can now engage, in effect, in a critical discussion about whether A is a reason for B by examining our shared background commitments and highlight the aspects that we take to be relevant to us having reasons for $A \to B$. Similarly, what matters for our ability to engage in a critical debate about whether A together with B is a reason against C is not whether either one of us accepts $\neg(A \wedge B \wedge C)$. Rather, what matters is whether we treat it as undeniable, that is, whether we take our shared background commitments to provide reasons for $\neg(A \wedge B \wedge C)$. I will do so, and you won't.

This is explicitation by implication. It must not be confused with the explicitation *of* implications that we mentioned in the previous chapter, where what is implied by a set of premises is added to these premises as a further premise. Explicitation *by* implication, by contrast, does not involve turning a conclusion into a premise. Rather, in explicitation by implication, we start with a particular reason relation. And according to this conception of explicitation, what it takes to make explicit a particular reason relation is to have a sentence that is undeniable just in case the reason relation indeed holds. That is, the reason relation holds if and only if the sentence is implied by the background commitments, against which we want to assess the particular reason relation at issue. We can summarize the general format of this kind of explicitation of reason relations in the following definition:

Definition 4 (Making Reason Relations Explicit by Implication)**.**
(For) The sentence ϕ makes explicit that Θ is a reason for Λ if and only if, for all Γ and Δ, we have $\Gamma \mathrel{|\!\sim} \phi, \Delta$ just in case $\Gamma, \Theta \mathrel{|\!\sim} \Lambda, \Delta$.
(Against) The sentence ϕ makes explicit that Θ is a reason against Λ if and only if, for all Γ and Δ, we have $\Gamma \mathrel{|\!\sim} \phi, \Delta$ just in case $\Gamma, \Theta, \Lambda \mathrel{|\!\sim} \Delta$.

It is easy to see that NMMS introduces logical vocabulary that allows us to make explicit by implication any reason relations among finite sets. To see this, let's write $\bigwedge\{x_1, ..., x_m\}$ for $x_1 \wedge ... \wedge x_m$ and $\bigvee\{x_1, ..., x_m\}$ for $x_1 \vee ... \vee x_m$. Now, by DD, AA, and SS above, for any pair of non-empty sets Θ and Λ, it is easy to see that $\Gamma, \Theta \mathrel{|\!\sim} \Lambda, \Delta$ if and only if $\Gamma \mathrel{|\!\sim} \bigwedge\Theta \to \bigvee\Lambda, \Delta$. So, for any (finite) Θ and Λ, the sentence $\bigwedge\Theta \to \bigvee\Lambda$ makes explicit that Θ is a reason for Λ, in the sense of (For). Similarly, by AA and II, we have $\Gamma, \Theta, \Lambda \mathrel{|\!\sim} \Delta$ if and only if $\Gamma \mathrel{|\!\sim} \neg\bigwedge\Theta \cup \Lambda, \Delta$. Hence, for any finite Θ and Λ, the sentence $\neg\bigwedge\Theta \cup \Lambda$ makes explicit that Θ is a reason against Λ, in the sense of (Against). We can summarize this result in the following fact.

Fact 5. *For any two finite sets of sentences, from the extended language $\mathfrak{L}$ for any base $\mathfrak{B}$, the logically extended lexicon of that base includes sentences that make explicit by implication when the relation of being a reason for and when the relation of being a reason against holds between the two sets. (Appendix, Proposition 28)*

This fact is one sense in which the vocabulary introduced with the help of NMMS is universally explicative of reason relations. For any autonomous base vocabulary, once we extend it to include logical vocabulary, we can make explicit by implication every reason relation among arbitrary finite sets of sentences.

We have now reached our aim of showing that the logical vocabulary of NMMS is universally LX. As shown in the previous section, the logical vocabulary of NMMS can be elaborated from any autonomous base vocabulary. And as just shown, it allows its users to make explicit arbitrary reason relations among sentences of any autonomous base vocabulary and its logical extension. This means that defining logical expressions for a given base vocabulary by the rules of NMMS meets the desiderata of logical expressivism.

We can now see what it means to adopt logical expressivism and to understand logic as the discipline that studies the vocabulary that universally allows us to make explicit what is a reason for and against what. If we use "conceptual self-consciousness" for the ability to be explicit about the exercises of one's own conceptual abilities and this means to make explicit the reason relations that constrain these exercises of conceptual abilities, then logic is the discipline that studies the vocabulary that allows us to manifest conceptual self-consciousness. Logic thus understood studies the tools that allow us to engage in critical reflection and discussion of the appreciation of reason relations of any kind whatsoever, including open reason relations. This is the conception of logic that animates logical expressivism, and the theory that we have developed in this chapter illustrates how one can formulate a logical theory based on this conception of logic. Thus, we have now put on the table a formally rigorous proposal for how to understand logical expressivism, and how to do logic in a way that is informed and guided by logical expressivism.

3.2.3 Explicitation by Sequents

Besides explicitation by implication, there is also a second and related way to spell out what it means to make reason relations explicit. The idea is that logical vocabulary makes reason relations explicit in the sense that any implication that involves logically complex sentences as premises or conclusions holds only if and because some definite implications among

only atomic sentences hold. In this subsection, we explain and explore this way to understand the idea that logical vocabulary makes reason relations explicit.

The second way to understand the idea that logical vocabulary makes explicit reason relations is to understand it as saying that all reason relations among logically complex sentences reflect or express reason relations among atomic sentences in a unique way. According to this way of understanding the idea, it is not a sentence that is implied by some premises that makes explicit reason relations. Rather, it is an entire implication that makes explicit particular reason relations among the base vocabulary. The idea is that every derivable sequent is, as it were, a projection of reason relations among sentences of the base vocabulary. On this way of thinking, for instance, the implication $p \mathrel{|\!\sim} q \rightarrow r, s$ makes explicit that the base consequence relation includes the implication $p, q \mathrel{|\!\sim} r, s$. The former implication is a projection of the latter; and note that a sequent featuring logically complex sentences might, in this sense, represent a collection of implications in the base. Let's call this "Explicitation by Sequents."

Definition 6 (Making Base Reason Relations Explicit by One Sequent)**.** Let $\texttt{AtomicImp} \subseteq \mathrel{|\!\underset{\mathfrak{B}}{\sim}}$ be a set of implications in the base consequence relation. Then the sequent $\Gamma \succ \Delta$ makes explicit (by one sequent) the collection of reason relations $\texttt{AtomicImp} = \{\Theta_1 \mathrel{|\!\underset{\mathfrak{B}}{\sim}} \Lambda_1, \ldots, \Theta_n \mathrel{|\!\underset{\mathfrak{B}}{\sim}} \Lambda_n\}$ if and only if ($\Gamma \mathrel{|\!\sim} \Delta$ just in case $\Theta_1 \mathrel{|\!\underset{\mathfrak{B}}{\sim}} \Lambda_1$, and ..., and $\Theta_n \mathrel{|\!\underset{\mathfrak{B}}{\sim}} \Lambda_n$).

The logical vocabulary of NMMS is such that for every implication in the extended consequence relation, there is a unique set of implications in the base consequence relation that it makes explicit in this sense. For the invertibility of the rules of NMMS ensures that the following holds:

Theorem 7 (*Projection*)**.** *For any sequent* $\Gamma \succ \Delta$, *with* $\Gamma \cup \Delta \subseteq \mathfrak{L}$, *there is a unique set* `AtomicImp` *of base vocabulary sequents such that* $\Gamma \succ \Delta$ *is derivable in* $\mathsf{NMMS}_{\mathfrak{B}}$ *just in case* $\texttt{AtomicImp} \subseteq \mathrel{|\!\underset{\mathfrak{B}}{\sim}}$. *(Appendix, Proposition 29)*

This means that every implication among sentences in the logically extended vocabulary makes explicit (by one sequent) a collection of reason relations in the base vocabulary. To put it differently, every sequent featuring logically complex vocabulary reflects a unique set of sequents featuring only atomic sentences: the complex sequent holds just in case all the atomic sequents hold.

It follows that for any set, `ComplexImp`, of sequents in the logically extended language, there is a unique set, `AtomicImp`, of sequents in

the nonlogical base vocabulary such that all the complex sequents in `ComplexImp` hold if and only if all the base sequents in `AtomicImp` hold. Moreover, this relation does not depend on what other sequents hold in the base or the extended vocabulary. Rather, this relation holds among the sets `ComplexImp` and `AtomicImp` for all base vocabularies and extensions. Hence, the set of complex sequents `ComplexImp` makes explicit (by sequents) the set of base sequents `AtomicImp` independently of what we choose as our base vocabulary. That Projection holds means that the logical vocabulary of NMMS can make arbitrary nonlogical reason relations explicit by sequents, in this sense. This is a second sense, beside explicitation by implication, in which the logical vocabulary of NMMS is universally explicative of reason relations. We take these two senses of being universally explicative—namely the sense of making arbitrary reason relations explicit by implication and by sequents—to be the two senses of being universally explicative that are central to logical expressivism.

We can now ask whether the converse is also true, that is, whether every set of implications in the base vocabulary is represented by a single implication in the logically extended consequence relation. Unfortunately, such a strong representation relation does not hold in NMMS. There are collections of base vocabulary implications that cannot be made explicit by a single sequent in the extended consequence relation. To see why, suppose that $p \mathrel{|\!\sim} q$ and $r \mathrel{|\!\sim} s$ but $p, r \not\mathrel{|\!\sim} q, s$. We can make the first two implications explicit by $\mathrel{|\!\sim} p \to q$ and $\mathrel{|\!\sim} r \to s$. But we cannot combine these two implications into a single implication. For, in order to derive $\succ (p \to q) \wedge (r \to s)$, we would need not only $\succ p \to q$ and $\succ r \to s$ but also $\succ p \to q, r \to s$. And the derivation of the latter sequent would require that $p, r \succ q, s$ is derivable, which it is not because $p, r \not\mathrel{|\!\sim} q, s$.

We could avoid such cases by requiring that base consequence relations are closed under the following rule, which we may call "Mingle-Mix": if $\Gamma \mathrel{|\!\sim} \Delta$ and $\Theta \mathrel{|\!\sim} \Lambda$, then $\Gamma, \Theta \mathrel{|\!\sim} \Delta, \Lambda$. If base consequence relations must obey Mingle-Mix, then it cannot happen that $p \mathrel{|\!\sim} q$ and $r \mathrel{|\!\sim} s$ but $p, r \not\mathrel{|\!\sim} q, s$. Unfortunately, Mingle-Mix is not plausible in a nonmonotonic setting. It seems plausible, for instance, that "The patient has symptom X" implies "The patient has disease XX" and "The patient has symptom Y" implies "The patient has disease YY"; but "The patient has symptoms X and Y" does not imply "The patient has disease XX or YY" (because the combination perhaps strongly suggests that the patient has disease ZZ). However, if we want to allow that Mingle-Mix fails in material consequence relations, then there are collections of implications in base consequence relations that cannot be made explicit by one single sequent in NMMS.

While this is a limitation of NMMS as formulated above, we can tweak our theory in such a way that our logical vocabulary is explicative of any

collections of reason relations by one sequent. To do so, we must reject Contraction, that is, the principle that $\Gamma, \phi \mathrel{|\!\!\sim} \Delta$ just in case $\Gamma, \phi, \phi \mathrel{|\!\!\sim} \Delta$ and that $\Gamma \mathrel{|\!\!\sim} \phi, \Delta$ just in case $\Gamma \mathrel{|\!\!\sim} \phi, \phi, \Delta$. Hence, we must work with multi-sets of sentences instead of sets of sentences. If we do that, we get the desired result if we change our sequent rules to the more familiar Ketonen-style rules, which are the following (marking the tweaked connectives by adding a bar above them):

$$\frac{\Gamma \succ \Delta, A \qquad B, \Gamma \succ \Delta}{\Gamma, A \bar{\rightarrow} B \succ \Delta}\,[\mathrm{L}\bar{\rightarrow}] \qquad \frac{\Gamma, A \succ B, \Delta}{\Gamma \succ A \bar{\rightarrow} B, \Delta}\,[\mathrm{R}\bar{\rightarrow}]$$

$$\frac{\Gamma, A, B \succ \Delta}{\Gamma, A \bar{\wedge} B \succ \Delta}\,[\mathrm{L}\bar{\wedge}] \qquad \frac{\Gamma \succ \Delta, A \qquad \Gamma \succ \Delta, B}{\Gamma \succ \Delta, A \bar{\wedge} B}\,[\mathrm{R}\bar{\wedge}]$$

$$\frac{\Gamma, A \succ \Delta \qquad \Gamma, B \succ \Delta}{\Gamma, A \bar{\vee} B \succ \Delta}\,[\mathrm{L}\bar{\vee}] \qquad \frac{\Gamma \succ \Delta, A, B}{\Gamma \succ \Delta, A \bar{\vee} B}\,[\mathrm{R}\bar{\vee}]$$

The negation rules and the rules with just one top sequent remain the same as in NMMS. We call the sequent calculus that results from this change $\mathsf{NMMS}^{\backslash ctr}$. To see that the logical vocabulary of $\mathsf{NMMS}^{\backslash ctr}$ can make explicit by implication arbitrary particular reason relations, it suffices to note that DD, II, AA, and SS all hold in $\mathsf{NMMS}^{\backslash ctr}$. Hence, the logical vocabulary of $\mathsf{NMMS}^{\backslash ctr}$ can make individual reason relations explicit by implications in the same way as the vocabulary of NMMS. However, in $\mathsf{NMMS}^{\backslash ctr}$ we can combine the results of such explicitation into a single sequent. For example, because of [R$\bar{\wedge}$], if $\succ p \bar{\rightarrow} q$ and $\succ r \bar{\rightarrow} s$ are both derivable, then so is $\succ (p \bar{\rightarrow} q) \bar{\wedge} (r \bar{\rightarrow} s)$, and vice versa. Thus, the limitation regarding explicitation of collections of reason relations by sequents that arose for NMMS does not arise for $\mathsf{NMMS}^{\backslash ctr}$. And in general, the logical vocabulary of $\mathsf{NMMS}^{\backslash ctr}$ is such that, for any (finite) collection of reason relations among (finite) sets, there is a sequent that makes exactly those reason relations explicit.[20]

Proposition 8. *For any finite set,* `AtomicImp`*, of sequents among atomic sentences, there is a sequent* $\Gamma \succ \Delta$*, with* $\Gamma \cup \Delta \subseteq \mathfrak{L}$*, such that* $\Gamma \succ \Delta$ *is derivable in* $\mathsf{NMMS}^{\backslash ctr}_{\mathfrak{B}}$ *just in case* $\texttt{AtomicImp} \subseteq \mathrel{|\!\!\sim}_{\mathfrak{B}}$.

This result allows us to formulate a representation theorem for $\mathsf{NMMS}^{\backslash ctr}$. For we can now see that in $\mathsf{NMMS}^{\backslash ctr}$ every finite set, `AtomicImp`, of implications in the base consequence relation is represented by a single sequent in the logically extended consequence relation, and every implication in the logically extended consequence relation represents

a unique set of implications in the base consequence relation. Moreover, there is a unique set, `ExtImp`, of sequents in the logically extended consequence relation that represent a given set, `AtomicImp`, of implications in the base consequence relation. So, for any sequent in `ExtImp` it holds just in case all the sequents in `AtomicImp` hold. This yields the following representation theorem:

Theorem 9 (*Representation*). *For every finite set,* `AtomicImp`*, of sequents in the base vocabulary of* $\mathfrak{L}_{\mathfrak{B}}$ *there is a set of sequents,* `ExtImp`*, in the extended vocabulary of* $\mathfrak{L}$ *introduced by* $\mathsf{NMMS}^{\backslash ctr}_{\mathfrak{B}}$*, such that, for every sequent in* `ExtImp`*, it holds if and only if all the sequents in* `AtomicImp` *hold.*

This theorem tells us that the logical vocabulary introduced by $\mathsf{NMMS}^{\backslash ctr}$ is explicative of reason relations in the sense that not only does every sequent featuring logical vocabulary reflect a collection of reason relations in the base consequence relation but also every finite collection of reason relations in the base vocabulary can be represented in a single sequent that features logical vocabulary.

The price one has to pay in order to have logical vocabulary that is universally explicative in the stronger sense of making arbitrary finite collections of reason relations explicit and, hence, underwriting the representation theorem just stated is that one must reject contraction and, thus, allow that whether or not some premises imply some conclusions may depend on how many times a premise or conclusion occurs in the implication. We will mostly work with NMMS and assume contraction (by working with sets), thus being content with the explicative power of Projection without Representation. We will, however, occasionally discuss the option of rejecting contraction and use $\mathsf{NMMS}^{\backslash ctr}$ instead of NMMS.

Moreover, we should acknowledge that there is another sense in which the logical vocabulary of NMMS may be thought to fail to make explicit reason relations. For, one may distinguish between weak and strong representation of the reason-for relation by saying that ϕ weakly represents that Θ is a reason for Λ if and only if, for all Γ and Δ, we have $\Gamma \mathrel{|\!\!\sim} \phi, \Delta$ just in case $\Gamma, \Theta \mathrel{|\!\!\sim} \Lambda, \Delta$. But ϕ strongly represents that Θ is a reason for Λ if and only if, for all Γ and Δ, we have $\Gamma \mathrel{|\!\!\sim} \phi, \Delta$ just in case $\Gamma, \Theta \mathrel{|\!\!\sim} \Lambda, \Delta$ and, also, $\Gamma \mathrel{|\!\!\sim} \neg\phi, \Delta$ just in case $\Gamma, \Theta \not\mathrel{|\!\!\sim} \Lambda, \Delta$. Given this terminology, the conditional of NMMS and $\mathsf{NMMS}^{\backslash ctr}$ represents the reason-for relation only weakly and not strongly. For, it is not the case that $\Gamma \mathrel{|\!\!\sim} \neg(\bigwedge \Theta \to \bigvee \Lambda), \Delta$ if and only if $\Gamma, \Theta \not\mathrel{|\!\!\sim} \Lambda, \Delta$. In an analogous way, the negation of conjunctions only weakly represents the reason-against relation. For, it is not the case that $\Gamma \mathrel{|\!\!\sim} \neg\neg \bigwedge \Theta \cup \Lambda, \Delta$ if and only if $\Gamma, \Theta, \Lambda \not\mathrel{|\!\!\sim} \Delta$. For the purposes of our project in this book, two quick remarks about this issue must suffice. (i) The idea of strongly representing the reason-for relation is fraught with

difficulties. The idea obviously requires, for instance, that, for every set and every conditional, the set implies either the conditional or its negation. But it seems wrong to think that we always either have reason for an arbitrary conditional or for its negation. (ii) The issue is related in interesting ways to the topic of validity predicates and the so-called v-Curry Paradox, and we think that the issue is best discussed in that context, which is different from our present context (see Beall and Murzi, 2013).[21] So, we set aside the idea of strongly representing reason relations.

This concludes our discussion of how logical vocabulary makes explicit reason relations. Let us take stock. We saw that the logical vocabulary of NMMS can make explicit arbitrary instances of the relation of being a reason-for and being a reason-against, in the sense of explicitation by implication. The logical vocabulary of NMMS can do this for any base reason relations between sets of sentences whatsoever, including nonmonotonic and nontransitive reason relations. In this sense, the logical vocabulary is universally explicative of reason relations. Moreover, the logical vocabulary of NMMS is such that every implication in the extended vocabulary reflects a unique set of implications in the base vocabulary. This is a second sense in which the logical vocabulary of NMMS makes reason relations explicit. However, the logical vocabulary of NMMS cannot represent arbitrary collections of implications in the base vocabulary in a single implication in the logically extended vocabulary. We can move to logical vocabulary that allows us to do this by rejecting Contraction and changing our sequent calculus to $\mathsf{NMMS}^{\backslash ctr}$. This yields an interesting representation theorem for the logical vocabulary of $\mathsf{NMMS}^{\backslash ctr}$. Overall, we have seen in the previous section that the logical vocabulary of NMMS can be elaborated from any autonomous base vocabulary. And we have seen in this section that the logical vocabulary of NMMS can make explicit arbitrary reason relations of any autonomous base vocabulary. Taking these results together, we have shown that the logical vocabulary of NMMS is universally LX. It is the kind of vocabulary that can play the role that is the essential and characteristic role of logical vocabulary, according to logical expressivism.

3.3 Making Local Structure Explicit

We have repeatedly emphasized the importance of allowing for failures of structural principles, especially failures of monotonicity and transitivity. Unless a theory allows for such failures, it cannot codify open reason relations. And if logical expressivism is right in claiming that the characteristic function of logical vocabulary is to make explicit reason relations in general, including open reason relations, then any such theory cannot provide an account of what logic makes explicit. Such a theory

might correctly specify the narrowly logical part of a consequence relation, but it cannot account for the essential explicitating function of logical vocabulary that shows up in the capacity of logically complex sentences that are not logically true to codify material reason relations.

Someone might acknowledge this importance of allowing for failures of structural principles while also insisting that structural principles hold in some areas of any consequence relation worth theorizing and that it is of crucial importance to our practice of giving and asking for reasons what these areas are. We agree. After all, we sometimes seem to tacitly assume that the reason relations that govern a particular discourse are transitive and monotonic, and this seems to happen in particular in very sophisticated and successful kinds of discourse, like those in mathematics, some parts of physics, and the like. Moreover, sometimes we try to exploit only the reason relations of classical logic and to codify all other reason relations in the form of explicit definitions, which we then use as premises. In light of such phenomena, it would be unsatisfying merely to be told that the usual structural rules are invalid and should be rejected. In this section, we show how we can move beyond such an unsatisfying position. In particular, we show how we can theorize structural features that hold locally in particular regions of consequence relations, and how we can introduce logical vocabulary that makes such local structural features explicit in the object language.

3.3.1 Regions of Monotonicity

In general, we hold that a good implication, $\Gamma \mathrel{|\!\sim} \Delta$, may be defeated by the addition of a further premise, such that $\Gamma, A \not\mathrel{|\!\sim} \Delta$. And we hold that this can also happen when we add further conclusions, such that it might be that $\Gamma \not\mathrel{|\!\sim} B, \Delta$. However, there will typically be a wide range of additions to an implication that do not defeat the implication. A pair of sets of sentences $\langle X, Y\rangle$ belongs to this range of non-defeating additions to $\Gamma \mathrel{|\!\sim} \Delta$ if and only if $\Gamma, X \mathrel{|\!\sim} Y, \Delta$, that is, just in case adding the first member to the premises and the second member to the conclusions yields another good implication. We call the set of such pairs of non-defeating additions the "range of subjunctive robustness" of an implication.[22]

Definition 10 (Range of subjunctive robustness, $\mathtt{RSR}\,\langle\Gamma, \Delta\rangle$)**.** Given a set of premises, Γ, and a set of conclusions, Δ, the range of subjunctive robustness of $\langle\Gamma, \Delta\rangle \in \mathcal{P}(\mathfrak{L}) \times \mathcal{P}(\mathfrak{L})$, written $\mathtt{RSR}\,\langle\Gamma, \Delta\rangle$, is the set of pairs, $\langle X, Y\rangle$, such that $\Gamma, X \mathrel{|\!\sim} Y, \Delta$; that is, $\mathtt{RSR}\,\langle\Gamma, \Delta\rangle = \{\langle X, Y\rangle \in \mathcal{P}(\mathfrak{L}) \times \mathcal{P}(\mathfrak{L}) : \Gamma, X \mathrel{|\!\sim} Y, \Delta\}$.

We can add any premises and conclusions to a good implication without defeating the implication, as long as the addition is in the implications range of subjunctive robustness. If we want to be explicit about an implication's range of subjunctive robustness, we can write $\Gamma \mid\!\sim^{\uparrow R} \Delta$ for $\forall \langle X, Y\rangle \in R$ $(\Gamma, X \mid\!\sim Y, \Delta)$. Then $\Gamma \mid\!\sim^{\uparrow R} \Delta$ if and only if $R \subseteq \mathtt{RSR}\langle \Gamma, \Delta\rangle$, and $\mathtt{RSR}\langle \Gamma, \Delta\rangle$ is the largest set, Z, for which $\Gamma \mid\!\sim^{\uparrow Z} \Delta$.[23] The larger the range of subjunctive robustness of an implication is, the more it approximates a sequent that holds persistently, that is, the more it approximates a sequent where monotonicity holds locally.

When the range of subjunctive robustness of an implication is maximal it includes all pairs of sets of sentences. If this is the case for $\Gamma \mid\!\sim \Delta$, then any application of [Weakening] yields another good implication, that is, for any $\langle X, Y\rangle \in \mathcal{P}(\mathfrak{L}) \times \mathcal{P}(\mathfrak{L})$ we have $\Gamma, X \mid\!\sim Y, \Delta$. We could write this as $\Gamma \mid\!\sim^{\uparrow \mathcal{P}(\mathfrak{L}) \times \mathcal{P}(\mathfrak{L})} \Delta$, but to avoid clutter we will write it simply thus: $\Gamma \mid\!\sim^{\uparrow} \Delta$. Hence, monotonicity holds locally at the sequent $\Gamma \mid\!\sim \Delta$ just in case $\Gamma \mid\!\sim^{\uparrow} \Delta$. Given that we think that monotonic consequence is an important and noteworthy part of consequence, we can now express this by saying that we are not only interested in $\mid\!\sim$ but also in $\mid\!\sim^{\uparrow}$, that is, not only in all good implications but also those implications at which monotonicity holds locally.

According to the conception of logic formulated in logical expressivism, logical vocabulary should ideally allow us to make explicit when monotonicity holds locally at an implication. With this interest in mind, the following facts about our sequent rules are noteworthy: Firstly, if the range of subjunctive robustness of a sequent includes all possible additions of atomic sentences, then it includes all possible additions of any sentences whatsoever. Or, in more formal terms:

Fact 11. *$X, \Gamma \succ \Delta, Y$ is derivable for all $\langle X, Y\rangle \in \mathcal{P}(\mathfrak{L}_{\mathfrak{B}}) \times \mathcal{P}(\mathfrak{L}_{\mathfrak{B}})$ if and only if $Z, \Gamma \succ \Delta, U$ is derivable for all $\langle Z, U\rangle \in \mathcal{P}(\mathfrak{L}) \times \mathcal{P}(\mathfrak{L})$; that is, $\mathcal{P}(\mathfrak{L}_{\mathfrak{B}}) \times \mathcal{P}(\mathfrak{L}_{\mathfrak{B}}) \subseteq \mathtt{RSR}\langle \Gamma, \Delta\rangle$ if and only if $\mathcal{P}(\mathfrak{L}) \times \mathcal{P}(\mathfrak{L}) \subseteq \mathtt{RSR}\langle \Gamma, \Delta\rangle$. (Appendix, Proposition 30)*

In other words, if we can weaken an implication with arbitrary sets of atomic sentences, then monotonicity holds locally at that implication. Hence, if we can keep track of when we can weaken a sequent with arbitrary sets of atomic sentences, then we can keep track of where monotonicity holds locally.

The second noteworthy fact is that our sequent rules are such that if all the top sequents of an application of a rule hold persistently, then the bottom sequent also holds persistently:

Fact 12. *If all the top sequents of an application of any rule of* NMMS *(or* NMMS$^{\backslash ctr}$*)*[24] *are good implications at which monotonicity holds locally, then the bottom sequent is a good implication at which monotonicity holds locally, that is, if for every top sequent* $\Gamma \succ \Delta$ *of an application of a rule of* NMMS *(or* NMMS$^{\backslash ctr}$*)* $\Gamma \mathrel{|\!\sim^{\uparrow}} \Delta$ *holds, then, for the bottom sequent* $\Theta \succ \Sigma$*, we have* $\Theta \mathrel{|\!\sim^{\uparrow}} \Sigma$*. (Appendix, Proposition 31)*

Given these two facts, we can introduce a new kind of sequent arrow, $\succ^{\uparrow}$, such that $\Gamma \succ^{\uparrow} \Delta$ is an axiom of our sequent calculus just in case for all $\langle X, Y\rangle \in \mathcal{P}(\mathfrak{L}_{\mathfrak{B}}) \times \mathcal{P}(\mathfrak{L}_{\mathfrak{B}})$ we have $X, \Gamma \mathrel{|\!\underset{\mathfrak{B}}{\sim}} \Delta, Y$. Moreover, we now stipulate that all the rules of NMMS apply if all the top sequents are of the new kind, and then a bottom sequent of the new kind can be derived. For example, if our base obeys Containment, we will have $p \succ^{\uparrow} p$. And we can then apply the rule [R$\neg$] to derive the sequent $\succ^{\uparrow} p, \neg p$. Similarly, [R$\vee$] now lets us derive $\succ^{\uparrow} p \vee \neg p$. By Fact 11, we know that $p \mathrel{|\!\sim^{\uparrow}} p$. And using Fact 12 we can infer that $\mathrel{|\!\sim^{\uparrow}} p \vee \neg p$. Hence, the sentence $p \vee \neg p$ follows monotonically—or persistently—from every set of premises. If we call the sequent calculus with the changes just specified NMMS$^{\uparrow}$, we can formulate this insight in general terms as follows:

Proposition 13. *The sequent* $\Gamma \succ^{\uparrow} \Delta$ *is derivable in* NMMS$^{\uparrow}$ *if and only if monotonicity holds locally at* $\Gamma \mathrel{|\!\sim} \Delta$*, that is,* Γ *implies* Δ *persistently. (Appendix, Proposition 32)*

Note that the monotonic part of our consequence relation may include implications that are not in the narrowly logical part of the consequence relation. We may, for example, include the following implication in our base consequence relation: "PP is a parent of CC." $\mathrel{|\!\sim}$ "CC is a child of PP." Moreover, we may want to include arbitrary applications of [Weakening] to this implication with sentences from our base lexicon in our base consequence relation. If we do that, this implication will hold monotonically in our logically extended consequence relation. However, the implication is not part of the narrowly logical part of our extended consequence relation because its derivation requires axioms that are not instances of Containment.

As logical expressivists, we don't just want a concept of implications at which monotonicity holds in our metalanguage. We would ideally also like to be able to make this fact explicit in our object language. More specifically, we would like to be able to make explicit, in the object language, that something is not just implied but that it is implied persistently. Persistence means the implication holds no matter what the additional context might be, that is, come what may. That is one clear sense

of "necessarily." If our premise is that PP is a parent of CC, then this does not just imply that CC is a child of PP. Rather, it implies that necessarily CC is a child of PP; it implies that, relative to our premise, CC cannot fail to be a child of PP. Hence we may write the object-language expression that we want to add as "$\Box$"—like the necessity operator in standard modal logics, which we will call the "Monotonicity Box." The idea is to use this operator to mark implications that have a maximal range of subjunctive robustness. So that "PP is a parent of CC" implies: "$\Box$ CC is a child of PP."

We can therefore expand our language $\mathfrak{L}$ to $\mathfrak{L}^{\Box}$ by including sentences in which $\Box$ occurs as a one-place connective, in the usual way. And let us add the following rules to $\mathsf{NMMS}^{\uparrow}$ and call the resulting calculus $\mathsf{NMMS}^{\Box}$, where square brackets around an upward arrow indicate that the upward arrow is optional, that is, the bottom sequent is derivable with and also without the upward arrow.

$$\frac{\Gamma, A \succ \Delta}{\Gamma, \Box A \succ \Delta}\ [\mathrm{L}\Box] \qquad\qquad \frac{\Gamma \succ {}^{\uparrow}A, \Delta}{\Gamma \succ {}^{[\uparrow]}\Box A, \Delta}\ [\mathrm{R}\Box]$$

The left-rule for the box ensures that $\Box A$ implies everything that A implies. And the right-rule says that if A is implied in a sequent decorated with an upward arrow, then $\Box A$ is implied. Continuing our example from above, it is easy to see that $\succ \Box(p \vee \neg p)$ is derivable in $\mathsf{NMMS}^{\Box}$, given that the base obeys Containment. In general, the following holds:

Proposition 14. *The sequent* $\Gamma \succ \Box A, \Delta$ *is derivable in* $\mathsf{NMMS}^{\Box}$ *if and only if* $\Gamma \mathrel{|\!\sim}^{\uparrow} A, \Delta$. *(Appendix, Proposition 33)*

This result captures the sense in which the Monotonicity Box allows us to make explicit local areas in which monotonicity holds. For it allows us to make explicit when something is not just implied but persistently implied. We can thus not only theorize monotonic consequence, including material monotonic consequence, in our meta-language but we can make this important structural feature of reason relations explicit in our object language. Moreover, just like the results above, this holds universally, that is, it holds for any base vocabulary whose reason relations we may want to make explicit.

It is worth noting that while our Monotonicity Box behaves like a standard necessity operator in many ways, it also differs in significant ways from such standard necessity operators. Among its familiar behavior is that, given any base that obeys Containment, all of the following implications hold, using "$\Diamond$" to abbreviate "$\neg\Box\neg$": $\Box A \mathrel{|\!\sim} \Diamond A$, and $\neg\Box A \mathrel{|\!\sim} \Diamond\neg A$; moreover, (K) $\Box(A \to B) \mathrel{|\!\sim} (\Box A \to \Box B)$, and (4) $\mathrel{|\!\sim} \Box(A \to A)$, and (B) $\mathrel{|\!\sim} A \to \Box\Diamond A$, and (D) $\mathrel{|\!\sim} \Box A \to \Diamond A$, and (5) $\mathrel{|\!\sim} \Diamond A \to \Box\Diamond A$ all hold.

Although these results look familiar from the modal logic S5, there are important differences. The necessitation rule is not sound in $\mathsf{NMMS}^{\Box}$ because it can happen that $\succ A$ is derivable but $\succ \Box A$ isn't derivable (and it perhaps should not be derivable because $\mid\!\sim^{\uparrow} A$ doesn't hold). Furthermore, some results would seem very strange for familiar modal logics, such as that the following implications hold, given a base that obeys Containment: $\mid\!\sim A \rightarrow \Box A$ and $\neg\Box A \mid\!\sim \neg A$. Although these results might seem strange from the perspective of standard modal logics, they are what one should expect, given that the Monotonicity Box makes explicit local monotonicity. For example, $\mid\!\sim A \rightarrow \Box A$ makes explicit by implication that A monotonically implies A, which is ensured by Containment. And $\neg\Box A \mid\!\sim \neg A$ holds because $\mid\!\sim \neg A, \Box A$, which holds because $\mid\!\sim^{\uparrow} \neg A, A$, which is in turn ensured by Containment and our negation rules.

To sum up, we have shown how to introduce a modal operator that makes explicit where monotonicity holds locally, that is, the conclusion $\Box A$ is implied just in case the conclusion A is implied persistently. If our base consequence relation obeys Containment, then all theorems of classical logic will be persistent theorems of our extended consequence relation. So, if ϕ is a theorem of classical propositional logic, then we have $\mid\!\sim^{\uparrow}\phi$ and, hence, $\mid\!\sim \Box\phi$. Thus, we can make explicit by implication that the theorems of classical logic are persistently implied by any set of premises. Moreover, whenever the implication $\Gamma \mid\!\sim \Delta$ is indefeasible, we have $\mid\!\sim^{\uparrow} \bigwedge\Gamma \rightarrow \bigvee\Delta$ and, hence, $\mid\!\sim \Box(\bigwedge\Gamma \rightarrow \bigvee\Delta)$. So we can make explicit by implication any implication that holds indefeasibly—saying of it in the logically extended object language that it holds indefeasibly.

3.3.2 Classicality, Contraction, and the General Case

In the previous subsection, we showed how one can make explicit when monotonicity holds locally at an implication. The strategy that we used there can be generalized to many other cases. To illustrate this, let us look at classical logic again.

We have acknowledged above that there are discourses in which classical logic is of crucial importance and all other implications are wrapped in definitions of non-logical terms that one can use as premises. In this kind of discourse, critical reflection and discussion may require that we can make explicit implications that hold in classical logic. As already intimated above, these are all and only the implications that can be derived in NMMS from axioms that are instances of Containment. Those are the only implications in a base consequence relation that hold in classical propositional logic. Now, the analogue of Fact 12 above holds for classicality, namely:

Fact 15. *If all the top sequents of an application of any rule of* NMMS *hold in classical propositional logic, then so does the bottom sequent.*

In parallel to what we did above, we can introduce a new kind of sequent arrow, $\succ^{cl}$, and say that $\Gamma \succ^{cl} \Delta$ is an axiom of our new calculus just in case $\Gamma \cup \Delta \subseteq \mathfrak{L}_{\mathfrak{B}}$ and $\Gamma \cap \Delta \neq \emptyset$, that is, if it is an atomic instance of Containment. Moreover, we add the one-place operator $(\!|cl|\!)$ to our language, and we add the following rules to NMMS, thus yielding a calculus we call $\mathsf{NMMS}^{(\!|cl|\!)}$, where the square brackets indicate that the "$[cl]$"-mark in the bottom sequent is optional:

$$\frac{\Gamma, A \succ \Delta}{\Gamma, (\!|cl|\!)A \succ \Delta}\ [\mathrm{L}(\!|\mathrm{cl}|\!)] \qquad\qquad \frac{\Gamma \succ^{cl} A, \Delta}{\Gamma \succ^{[cl]} (\!|cl|\!)A, \Delta}\ [\mathrm{R}(\!|\mathrm{cl}|\!)]$$

It is then easy to prove the following result:

Proposition 16. *If all the sentences in* $\Gamma \cup \Delta \cup \{A\}$ *are in the language of classical propositional logic, then the sequent* $\Gamma \succ (\!|cl|\!)A, \Delta$ *is derivable in* $\mathsf{NMMS}^{(\!|cl|\!)}$ *if and only if* $\Gamma \vdash_{CL} A, \Delta$. *(Appendix, Proposition 34)*

We can thus make explicit the classicality of an implication in the same generic way in which we can make explicit that an implication holds monotonically. As already intimated above, there can be monotonic implications that are not classical implications. Moreover, there can be monotonic implications that are not transitive, whereas classical implications are closed under Cut. So, the Monotonicity Box and the classicality operator are genuinely distinct notions.

Let us now step back and consider the overall strategy at work here. The general recipe behind the operator that makes monotonicity explicit and the operator that makes classicality explicit is the following. (a) There is a feature some implications in our base vocabulary have. (b) This feature is preserved by the rules of NMMS, that is, if all the top sequents in a rule application have the feature, then so does the bottom sequent. (c) An implication in the logically extended language has the feature only if the corresponding sequent is derivable from just those axioms that have the feature. If (a)–(c) hold, we can introduce a new sequent arrow whose axioms are the base sequents with the feature at issue. We can read all the sequent rules as applying to the new kind of sequent arrow, and we can add rules that allow us to add a new operator on the left without further ado and allow us to add the operator on the right if the top sequent uses the new kind of sequent arrow. The result is an operator that makes explicit implications in the logically extended vocabulary that have the feature at issue.

This recipe can be applied in many cases. We have seen how it applies in the cases of monotonicity and classicality. It can also be applied to Contraction, which can fail in $\mathsf{NMMS}^{\backslash ctr}$, as explained above. So let us run through the general recipe again for this case. Let us use $\mathbb{F}$ to talk about the relevant feature, in our case Contraction. For the case of contraction, we can define our feature as follows: $\mathbb{F}(\Gamma \mathrel{|\!\sim} \Delta)$ just in case $\Gamma \mathrel{|\!\sim} \Delta$ and, for all $\Theta \succ \Sigma$, if $\Theta \succ \Sigma$ can be derived from $\Gamma \succ \Delta$ by just applications of Contraction, then $\Theta \mathrel{|\!\sim} \Sigma$.

(a) Some implications, $\Gamma \mathrel{|\!\sim}_{\mathfrak{B}} \Delta$ in our base are such that, for all $\Theta \succ \Sigma$, if $\Theta \succ \Sigma$ can be derived from $\Gamma \succ \Delta$ by just applications of Contraction, then $\Theta \mathrel{|\!\sim}_{\mathfrak{B}} \Sigma$. Those are the base implications such that $\mathbb{F}(\Gamma \mathrel{|\!\sim}_{\mathfrak{B}} \Delta)$.

(b) Given the rules of $\mathsf{NMMS}^{\backslash ctr}$, if all the top sequents of an application of a rule have feature $\mathbb{F}$, then the bottom sequent has feature $\mathbb{F}$. That is, if for all the top sequents it is admissible to apply the Contraction rule, then it is admissible to apply the Contraction rule to the bottom sequent.

(c) If some sequent derivable in $\mathsf{NMMS}^{\backslash ctr}$ has feature $\mathbb{F}$, then it is derivable from sequents that have feature $\mathbb{F}$. That is, all sequents for which the application of the Contraction rule is admissible, can be derived from sequents for which the Contraction rule is admissible.

Since (a)–(c) hold for the feature of Contraction being admissible, we can introduce a new sequent arrow, $\succ^{f}$, and say that $\Gamma \succ^{f} \Delta$ is an axiom just in case $\mathbb{F}(\Gamma \mathrel{|\!\sim}_{\mathfrak{B}} \Delta)$. We allow the rules of $\mathsf{NMMS}^{\backslash ctr}$ to apply when all top-sequents use the new kind of sequent arrow $\succ^{f}$. And we add the following rules:

$$\frac{\Gamma, A \succ \Delta}{\Gamma, (\!|f|\!)A \succ \Delta}\ [\mathrm{L}(\!|f|\!)] \qquad\qquad \frac{\Gamma \succ^{f} A, \Delta}{\Gamma \succ^{[f]} (\!|f|\!)A, \Delta}\ [\mathrm{R}(\!|f|\!)]$$

And we can then show that our new operator makes explicit the feature at issue, which is the admissibility of Contraction in the present case. We can formulate the result in general terms as follows.

Proposition 17. *The sequent* $\Gamma \succ (\!|f|\!)A, \Delta$ *is derivable in* $\mathsf{NMMS}^{\backslash ctr, (\!|f|\!)}$ *if and only if* $\mathbb{F}(\Gamma \mathrel{|\!\sim} A, \Delta)$.

For the particular case of the admissibility of Contraction, this result tells us that our new operator allows us to make explicit when an implication holds in such a way that Contraction can be safely applied.

3.3.3 Cautious Monotonicity and Cumulative Transitivity

With the general strategy from the previous subsection in hand, we may turn to two structural principles that are often deemed desirable in nonmonotonic logics, but which we have argued in the previous chapter are not desiderata for our project. These are Cautious Monotonicity (CM) and Cumulative Transitivity (CT), which can be formulated thus:

$$\frac{\Gamma \succ A \qquad \Gamma \succ \Delta}{\Gamma, A \succ \Delta}\ [\mathrm{CM}] \qquad\qquad \frac{\Gamma \succ A \qquad \Gamma, A \succ \Delta}{\Gamma \succ \Delta}\ [\mathrm{CT}]$$

As already intimated in the previous chapter, CM says that one can never lose consequences by making implications explicit, in the sense of adding them to one's premises. And CT says that one can never gain consequences by making implications explicit in this sense. Thus, these two structural principles together say that explicitation—in the sense of adding implications to one's premises—is inconsequential. We reject this thesis, and we hold that explicitation can be consequential.

Interestingly, local areas in which CM or CT hold cannot easily be made explicit in the way we can make explicit monotonicity, classicality, and the admissibility of contraction. It is instructive to see why this is so. Let us start with CM.

The problem with CM as a global constraint is that, in the context of a conditional that is adequate for logical expressivism, it implies a much stronger condition. Recall that, according to logical expressivism, the conditional must obey the Deduction-Detachment (DD) condition in order to perform its explicitation function adequately. That is, $\Gamma \mathrel{|\!\sim} A \to B, \Delta$ just in case $\Gamma, A \mathrel{|\!\sim} B, \Delta$. That is what it is for the conditional to make the metalinguistic implication turnstile explicit in the logically extended object language. Now, suppose that $\Gamma, B, C \mathrel{\underset{\mathfrak{B}}{|\!\sim}} D$ and $\Gamma \mathrel{\underset{\mathfrak{B}}{|\!\sim}} A$ but $\Gamma, B, C \mathrel{\underset{\mathfrak{B}}{\not|\!\sim}} A$. In the base consequence relation, CM does not require that $\Gamma, A, B, C \mathrel{\underset{\mathfrak{B}}{|\!\sim}} D$ and we can, hence, stipulate that $\Gamma, A, B, C \mathrel{\underset{\mathfrak{B}}{\not|\!\sim}} D$. For A is not implied by $\Gamma \cup \{B, C\}$; it is merely implied by a proper subset of them, namely Γ. However, it follows from DD that $\Gamma, B, C \mathrel{|\!\sim} D$ just in case $\Gamma \mathrel{|\!\sim} B \to (C \to D)$. We can now apply CM, and this yields $\Gamma, A \mathrel{|\!\sim} B \to (C \to D)$, which by DD entails that $\Gamma, A, B, C \mathrel{|\!\sim} D$. But that contradicts our stipulation above.

If we insist that our conditional obey DD, as logical expressivism suggests, and we insist that the logical extension of our base consequence relation is conservative, then the only way to ensure that CM holds is to enforce a stronger principle, which we call "Weakening with Implications of Subsets" (WIS):

$$\frac{\Gamma \succ A \qquad \Gamma, \Theta \succ \Delta}{\Gamma, \Theta, A \succ \Delta}\ [\text{WIS}]$$

This principle says that if some premise-set implies a conclusion, say $\Gamma \cup \Theta \mathrel{|\!\sim} D$, and a subset of this premise-set implies something, say $\Gamma \mathrel{|\!\sim} A$, then we can weaken the original inference by adding the implication of the subset of the premises as another premise, $\Gamma \cup \Theta, A \mathrel{|\!\sim} D$. Thus, this principle allows us to weaken any inference with any sentence that is implied by any subset of the premises of the inference.

If a base consequence relation is closed under WIS, then WIS and, hence, CM hold in the logically extended consequence relation. However, WIS is too strong to be plausible. For, we want to allow for situations like the following: "This is a chair" implies "You can sit on this." "This is a chair" and "This is a piece of art in an exhibition" implies "You are not allowed to touch this." However, if we combine the three claims as premises, "This is a chair" and "This is a piece of art in an exhibition" and "You can sit on this", they do not imply "You are not allowed to touch this." After all, there are exhibitions that invite the audience to participate by using the artwork.

To sum up, if one is willing to accept WIS as a constraint on base consequence relations, then our logical extensions obey WIS and, hence, CM. Given the constraints of logical expressivism that logical extensions of reason relations are conservative and the conditional satisfies DD, this is the only way to ensure that CM holds as a global structural principle.

If we don't want to restrict ourselves to bases that obey WIS, we might still want to make explicit when CM holds locally at an implication. However, an interesting question arises here, namely what we mean by "CM holds locally at an implication." One thing one can mean by this is that CM holds at an implication $\Gamma \mathrel{|\!\sim} A$ if and only if any implication whose premises are exactly Γ can be weakened, on the left side, with A.

Definition 18 (CM holding locally)**.** CM holds locally at an implication $\Gamma \mathrel{|\!\sim} A$ just in case $\Gamma \mathrel{|\!\sim} A$ and, for all B, if $\Gamma \mathrel{|\!\sim} B$, then $\Gamma, A \mathrel{|\!\sim} B$.

If we want to make explicit by implication when CM holds locally in this sense, we will try to introduce an operator, $(\!|cm|\!)$, such that $\Gamma \mathrel{|\!\sim} (\!|cm|\!)A$ if and only if CM holds locally at $\Gamma \mathrel{|\!\sim} A$. However, this idea is self-undermining in the following way. Suppose that CM holds locally at $\Gamma \mathrel{|\!\sim} A$ and $\Gamma \mathrel{|\!\sim} B$. Moreover, $\Gamma \mathrel{|\!\sim} C$ but $\Gamma, A, B \not\mathrel{|\!\sim} C$. Since CM holds at $\Gamma \mathrel{|\!\sim} B$, we can weaken the implication of C to get $\Gamma, B \mathrel{|\!\sim} C$. Given how we would like to make explicit when CM holds locally, we will also have $\Gamma \mathrel{|\!\sim} (\!|cm|\!)A$ and $\Gamma \mathrel{|\!\sim} (\!|cm|\!)B$. It follows from these two sequents that $\Gamma, B \mathrel{|\!\sim} (\!|cm|\!)A$. However, this is wrong; we have $\Gamma, B \mathrel{|\!\sim} C$ but $\Gamma, A, B \not\mathrel{|\!\sim} C$. So there is an implication of the premise-set $\Gamma \cup \{B\}$ that cannot be weakened with

A, contradicting what we want $\Gamma, B \mathrel{|\!\sim} (\!|cm|\!) A$ to mean. Hence, while $\Gamma \mathrel{|\!\sim} (\!|cm|\!) B$ holds if we consider just the implications that do not include our new operator $(\!|cm|\!)$, the very introduction of this operator makes it the case that $\Gamma \mathrel{|\!\sim} (\!|cm|\!) B$ no longer holds.

We call this kind of phenomenon an "expressive paradox." And what we mean by this is a situation in which something holds of an implication or a set of implications but when we add to our object language the resources to make explicit that this holds of the implication(s), it no longer holds. The situation is similar to the attempt to introduce a sentence whose utterance would allow one to make explicit that one is currently not uttering any sentence.

While there may be ways to make explicit when CM holds locally that do not give rise to expressive paradoxes, we shall not pursue this project any further here. We want to point out, however, that the existence of expressive paradoxes implies that the project of logical expressivism has limits. Some features of reason relations may be such that they cannot be made explicit without undermining them. And sometimes it might not be possible to make two things both explicit at the same time, so that we have to choose between the ability to make one explicit and the ability to make the other explicit. We are inclined to think that semantic paradoxes like the Liar Paradox and Curry's Paradox can be understood as expressive paradoxes and that seeing semantic paradoxes as particular instances of the wider class of expressive paradoxes might be fruitful. However, we will not pursue that idea in this work.

Finally, let us turn to Cumulative Transitivity (CT). There is no fully satisfying way to make explicit when CT holds in a region of our consequence relation. However, we can make explicit in which regions CT together with Monotonicity hold in the following way (see Hlobil, 2017).

Definition 19 (MOT-base-regions). An MOT-base-region is a subset of $\underset{\mathfrak{B}}{\mathrel{|\!\sim}}$ that is closed under [Weakening] and under CT, that is, a subset such that (i) if $\Gamma \underset{\mathfrak{B}}{\mathrel{|\!\sim}} A$ and $\Gamma, A \underset{\mathfrak{B}}{\mathrel{|\!\sim}} \Delta$, then $\Gamma \underset{\mathfrak{B}}{\mathrel{|\!\sim}} \Delta$, and (ii) if $\Gamma \underset{\mathfrak{B}}{\mathrel{|\!\sim}} A$, then $\Gamma, B \underset{\mathfrak{B}}{\mathrel{|\!\sim}} A$.

If $mot.n$ is the nth MOT-base-region in which we are interested, we introduce a sequent arrow $\succ^{mot.n}$ and let $\Gamma \succ^{mot.n} \Delta$ be an axiom of our new calculus NMMS^{mot} just in case $\Gamma \underset{\mathfrak{B}}{\mathrel{|\!\sim}} \Delta$ holds and is in the MOT-base-region $mot.n$. Each MOT-base-region will then have a logical extension, consisting of the sequents derivable from axioms that are all in the MOT-base-region. This logical extension of the MOT-base-region will also be closed under CT and [Weakening]. We can then introduce an operator that allows us to make explicit which implications belong to this logical extension, namely as follows:

$$\frac{\Gamma, A \succ^{mot.n} \Delta}{\Gamma, (\!|mot.n|\!)A \succ^{[mot.n]} \Delta}\ [\mathrm{L}(\!|\mathrm{mot.n}|\!)] \qquad \frac{\Gamma \succ^{mot.n} A, \Delta}{\Gamma \succ^{[mot.n]} (\!|mot.n|\!)A, \Delta}\ [\mathrm{R}(\!|\mathrm{mot.n}|\!)]$$

We can then show that $\Gamma \mathrel{|\!\sim} (\!|mot.n|\!)A, \Delta$ holds just in case $\Gamma \mathrel{|\!\sim} A, \Delta$ is in the logical extension of the MOT-base-region *mot.n*. Hence, if $\Gamma \mathrel{|\!\sim} (\!|mot.n|\!)A, \Delta$, then [Weakening] and CT are admissible for this implication if the second top premise of the CT application is also in the logical extension of the MOT-base-region *mot.n* (see Hlobil, 2017).

We thus have a way to keep track, in the object language, of regions of our consequence relations in which Monotonicity and Cumulative Transitivity hold, that is, regions that are closed under these rules. This approach could be refined and further developed. For our current purposes, however, it suffices to note that various structural features of consequence relations can be made explicit in the object language. The degree of difficulty and the amount of added complications vary between different structural features. When Monotonicity or classicality hold logically at an implication, we can make this explicit in a straightforward way by using the recipe described in the previous two subsections. This recipe works also when we reject Contraction as a global structural principle and want to make explicit when Contraction is nevertheless admissible for particular sequents. The cases of Cautious Monotonicity and Cumulative Transitivity give rise to more complications. We have thus sketched the beginning of an account of how and to what degree local structural features of consequence relations can be made explicit. What we have said is merely a start in this direction. It suffices, however, to illustrate in what kind of new logical vocabulary a logical expressivist might be interested and how such new logical vocabulary might be introduced. The logical expressivist is interested in making explicit reason relations and their local structural features. And in this section, we have seen how this interest gives rise to questions and projects that have not been addressed, as far as we know, by any extant logical theories.

3.4 Conclusion

The previous chapter introduced a stark distinction between two ways of thinking about the relations between the two topics of our title: logic and reasons. In the simplest terms, logicists about reasons understand reasons in terms of logic: good reasons are, in the end, logically good reasons. Expressivists about logic understand logic in terms of reasons: the defining job of logical vocabulary is to make explicit reason relations of implication and incompatibility. Expressivists must understand the reason relations that logical vocabulary codifies as settled elsewhere, before logic comes on the scene. We have sketched a pragmatics-first order of explanation,

according to which those reason relations among sentences are instituted by the norms governing their use in discursive practices of making claims and challenging and defending them. We talk about "reason relations" because being a reason is a relational property, and because reason relations merely constrain and do not determine what one, ultimately, has reason to accept or reject. Indeed, it is a bivalent relational property: reasons are essentially, and not just accidentally, either reasons for or reasons against—hence the two kinds of reason relations: implication and incompatibility. In the next chapter we begin to consider a semantics-first order of explanation, according to which the reason relations are settled by the relations of sentences to worldly states that make them true or false. Also in Chapter Four, we open the discussion of how to understand the relations between those normative pragmatic and representational semantic perspectives on the basic reason relations that expressivists understand as providing the raw materials that metalogical machinery then puts into explicit logical shape.

Since its Fregean origins, modern logic has been largely framed, shaped, and conducted within the confines of a logicist philosophical understanding of the enterprise. (We think that Frege himself was a rational expressivist, but the motivation of our project does not turn on that controversial hermeneutic claim.) The question we addressed in this chapter is how one can and ought to do logic differently, if one instead understands the enterprise philosophically in logical expressivist terms. What kind of logic is motivated by expressivism in the philosophy of logic? How can one best deploy what we have learned technically about logic during its development under logicist auspices, in the service of fulfilling the aspirations of logical expressivism? In subsequent chapters we ask a corresponding question about model-theoretic formal semantics—with special attention to the semantics of logical vocabulary.

In keeping with its logicism, the traditional philosophical understanding of logic focuses on pure logic: reason relations that hold in virtue of logic alone. In the sequent calculus, this means looking only at proof trees whose leaves are instances of Containment, that is, derivations from initial sequents of the form $\Gamma, A \succ A, \Delta$. By contrast, we think logic appears in its most characteristic and revealing guise when it is applied. In the sequent calculus, this means looking at proof trees whose leaves are sequents codifying substantial material implications—those whose goodness is underwritten not by a structural principle such as Containment but rather by the norms governing the use of the nonlogical sentences involved, or the way they represent the world as being.

We can also think of the distinction from the side of analysis rather than synthesis. (Perhaps it seems to be begging the question in favor of logical expressivism to think of things in terms of applying to an antecedent

field of material reason relations connective rules that we have stipulated.) Given a set of relations of implication and incompatibility among sets of sentences (including logically complex ones), and a distinction between logical and nonlogical vocabulary, reason relations that hold in virtue of logic alone (those that are due to the "logical form" of the sentences involved rather than their nonlogical content) can be picked out by using the Bolzano-Frege-Quine method of noting invariance under substitution. The logically good implications, are those good implications that remain good under uniform substitution of nonlogical for nonlogical vocabulary (and similarly for incompatibilities), including substitutions that include possible extensions or variations of the nonlogical vocabulary. Substitutionally picking out this subset of purely logical reason relations from an antecedently specified field of reason relations by seeing which ones hold *salva consequentia* only works if one can both tell good from bad implications without appeal to logically good implications, and can demarcate logical from nonlogical vocabulary. Logical expressivism as we have introduced it shows up as an option in the context of an order of explanation that assumes competent linguistic practitioners must have at least a rough and ready capacity to tell what is a reason for and against what—that is, to distinguish successful from unsuccessful rational defenses of and challenges to doxastic commitments (quite apart from whether that feature of the pragmatics is in addition understood as reflecting or instituting representational semantic properties of the sentences involved).

Against this background, logical expressivism then offers its distinctive criterion of demarcation of specifically logical vocabulary. Logical vocabulary is to be distinguished by its characteristic expressive role: making reason relations explicit. Here "making explicit" is formulating relations of implication and incompatibility as the content of declarative sentences, which are understood as what can both be used to defend or challenge other claims by offering reasons for or against them, and can itself be defended or challenged by other claims. In Chapter Two, where the expressivist understanding of the relation between logic and reasons was introduced, it was argued that implicit in this conception of the expressive task distinctive of logical vocabulary is a regulative ideal: logical vocabulary should be elaborated from and explicative of any and every base vocabulary. In our slogan, it should be *universally LX*. Given *any* set of reason relations on a sentential lexicon ("universally"), the rules for introducing logical vocabulary should make it possible to compute ("elaborate"—the "L" in "LX") the reason relations of the logically extended vocabulary entirely from the reason relations of the base vocabulary. Further, when used in accordance with the reason relations elaborated from those of the base vocabulary, the new, logically complex compounds of the original sentences should make it possible to say

explicitly ("explicate"—the "X" in "LX") what all those reasons relations are—both those of the base and those of the extended vocabulary. Saying, making explicit, in the relevant sense consists in putting into the form of declarative sentences—that is, in claimable form, as what can be asserted and denied, challenged and defended.

In this chapter, we have introduced the sentential logic that is our best candidate for satisfying this expressive ideal, and shown how it satisfies all the basic criteria of adequacy logical expressivism motivates. Along the way, doing that required substantially clarifying and making more concrete those criteria of adequacy. The metavocabulary we use to introduce logical vocabulary is Gentzen's sequent calculus. It is ideal for our purposes because it treats reason relations (in the form of sequents) as the objects that it manipulates and operates on. Its connective rules, individually and collectively, define functions from reason relations codified in a set of sequents relating sets of sentences of the base vocabulary to reason relations codified in the form of sequents relating sets of sentences of the logically extended vocabulary. The sequent calculus is accordingly explicitly designed, and perfectly suited, to perform the aspect of the function of logical vocabulary that consists in elaborating or computing one set of reason relations from another. It operationalizes the sense in which anyone who knows how to use the base vocabulary to make, challenge, and defend claims expressed by its sentences thereby knows how to do everything they need to know how to do in order to make, challenge, and defend claims expressed by all the sentences of the logically extended vocabulary. From an expressivist point of view, the sequent calculus is the metavocabulary that most perspicuously displays this essential aspect of the fundamental expressive task of logic. In Chapter Five we introduce implication-space semantics as the corresponding ideally perspicuous metavocabulary for computing the conceptual roles played by more complex sentences from those played by simpler ones—where, in accordance with semantic inferentialism, the conceptual roles of sentences are understood as the roles they play in reason relations.

The core of what expressivism in the philosophy of logic asks of logic is contained in the explicitation condition of the "universal LX-ness" formulation. That is the requirement that to qualify as logical, vocabulary must function to make explicit the reason relations of base vocabularies to which it is applied. The fundamental reason relations are implication or consequence, and incompatibility, corresponding to reasons for and reasons against. We began by formulating a precise sense in which familiar logical connectives can count as making such relations explicit, which we denominated as "explicitation by implication." Though they appear in slightly generalized form in the multisuccedent sequent-calculus formulation, explicitation by implication for the basic connectives

of sentential logic is epitomized in the two principles governing the core expressive connectives that codify implications and incompatibilities: the Deduction-Detachment (DD) condition on conditionals, which says that $\Gamma \mid\!\sim A \rightarrow B$ if and only if $\Gamma, A \mid\!\sim B$ and the Incoherence-Incompatibility (II) condition on negation, which says that $\Gamma \mid\!\sim \neg A$ if and only if $\Gamma, A \mid\!\sim$. (The auxiliary, aggregative connectives of conjunction and disjunction are governed by the corresponding Antecedent-Adjunction (AA) and Succedent-Summation (SS) conditions.) The DD condition ensures that a context implies the conditional $A \rightarrow B$ just in case, in that context, the antecedent of the conditional A implies its consequent B. That is one clear thing to mean by saying that the conditional codifies in the logically extended object language the implications that are specified by the turnstile in the sequent-calculus metalanguage. In that specific sense, the conditional says that an implication is good. That is explication by implication. Where DD holds, any implication can be explicitly expressed in the form of a conditional. Similar remarks hold of the other connective conditions.

Explicitation, the 'X' portion of "Universal LX-ness," in the sense of explicitation by implication, requires satisfaction of the four conditions DD, II, AA, and SS. The logic NMMS does so. Those conditions are all biconditionals. The metalogical tractability and expressive power of Gentzen's sequent calculus (its enabling the capacity to prove things about what things can be proven), however, depends on requiring that derivations using connective rules always proceed from the logically simpler to the logically more complex. That is, each step in a derivation, from the leaves to the root of a proof tree, involves adding connectives. Gentzen's student Ketonen saw that using the same distinction between derivability and admissibility that is exploited in Gentzen's Cut-eliminability "Hauptsatz" makes possible invertible rules—which permit the sort of biconditionals explicitation by implication requires. For even though derivations always move from the less to the more complex, it can still be arranged that whenever the sequent below the sequent-derivation line is derivable, so are the sequents above the line. Then the converse rule is admissible. That makes it possible to have invertible rules, which if properly chosen can satisfy DD, II, AA, and SS.

The Deduction-Detachment condition, the Incoherence-Incompatibility condition and the other conditions on connective definitions provide sufficient expressive power to guarantee the explicability by implication of any and all individual reason relations. But as we have arranged the logical extension of arbitrary base vocabularies by NMMS, there is also a kind of group or collective explication, which is a further kind of expressive completeness. For any NMMS implication relating sets of logically complex sentences, we can compute exactly what set of implications must hold in any base vocabulary for its logical extension to include that NMMS

implication. That is, each implication of one set of logically complex sentences by another codifies some particular set of implications in the base vocabulary. We can understand that as what it says about the reason relations of the base vocabulary. For the logical vocabulary of the non-contractive calculus $\mathsf{NMMS}^{\backslash ctr}$, the converse also holds: every set of implications in the base vocabulary can be made explicit by a single sequent in the logically extended vocabulary. The logic has the expressive power to make explicit in a single implication the fact that any set of implications holds in any base vocabulary. The clearly defined notions of explication by implication, projection, and representation of logical vocabulary relative to the base vocabularies to which it is applied make precise the idea of making reason relations explicit that is at the heart of logical expressivism. We show that NMMS satisfies that expressive ideal: it is LX for any and all base vocabularies.

In sequent calculi, connective rules define functions from reason relations to reason relations within the confines specified by the structural rules. The aspiration to expressive universality of logic with respect to reason relations—its being LX for any antecedent constellation of implications and incompatibilities—would be compromised insofar as the applicability, tractability, and expressive functionality of the connective rules depended on structural presuppositions. The stronger those presuppositions (Monotonicity and Transitivity as mixed-context Cut, or only Cautious Monotonicity and Cumulative Transitivity as shared-context Cut, etc.) the greater the infringement of the expressive ideal. We show how to relax those structural constraints—indeed, to eliminate them (though retaining Containment is shown to pay big dividends).[25] Doing that requires carefully tuning the connective rules so that they do not implicitly reimpose structures that have not been explicitly stipulated in the form of global structural rules. Since we hold, contrary to relevant logics, that licensing, for instance, the move from $\Gamma, A \mathrel{|\!\sim} \Delta$ to $\Gamma, A \wedge B \mathrel{|\!\sim} \Delta$ is a way to implicitly reimpose a kind of Monotonicity, relaxing the structural constraints in such a careful way requires mixing additive and multiplicative rules for conjunction and disjunction.[26]

In these ways, logical vocabulary can make explicit arbitrary reason relations, according to the theory developed in this chapter. That is, the logical vocabulary that we have introduced is universally explicative of reason relations, including open reason relations. Hence, we have shown how we can introduce logical vocabulary that is universally LX—that is, that can be elaborated from any base vocabulary and make explicit arbitrary reason relations. Crucially, this includes reason relations that are structurally open rather than closed in various senses. Even radically substructural reason relations can still be codified logically: not only those in which Monotonicity and Transitivity fail, but even consequence

and incompatibility relations that are hypernonmonotonic, in failing even Cautious Monotonicity.

We have also shown how vocabulary can be introduced that makes explicit where structural features hold locally in a consequence relation, even where those structural constraints do not hold globally. This includes local regions where Monotonicity, Classicality, and Contraction hold. This new vocabulary for marking local regions of sequents where reason relations satisfy various closure-structural conditions qualify as a kind of logical vocabulary by our expressivist criterion of demarcation. For instance, the intensional modal operator marking persistent (monotonic) sequents—a kind of necessity corresponding to a maximal range of subjunctive robustness of an implication—counts as a bit of logical vocabulary, since it makes explicit a feature of reason relations: that not only is an implication good, but so are all the implications related to it by having premise-sets and conclusion-sets that are supersets of the originals. By contrast, for instance, the set-theoretic epsilon does not qualify as a bit of logical vocabulary, since it does not make explicit any feature of reason relations as such. The new operators that we have introduced for codifying the structural properties of reason relations illustrate that logical expressivism gives rise to new interests and research questions, which extant logical theories have not raised or addressed.

Things had to be arranged so that all three of these kinds of expressive-explicative relations—explication by implication of individual sequents, the capacity to express collections of sequents with logically complex sequents, and the explicit marking with sentential operators of local structural features of sequents—continue to work as desired in the absence of global closure structure. That means that even if the base vocabulary is nonmonotonic (or even hypernonmonotonic, in failing Cautious Monotonicity), any set of base sequents is still expressible explicitly as a set of sequents in the logically extended vocabulary—and so, if we assume the bases are finite, in a single sequent relating sentences in the logically extended lexicon. For this to hold, it must be that the reason relations of the logically extended vocabulary are also open-structured. The extended vocabulary is a conservative extension of the base vocabulary, so at least those substructural or open-structured base sequents will remain so. But their open structure will also be reflected in the structure of the sequents that hold between sets of sentences in the logically extended lexicon. The expressive role of the logical connectives requires that the relations of consequence and incompatibility of the logically extended vocabulary do not satisfy closure-structural principles, so as to be able to make explicit the reason relations of base vocabularies that do not exhibit those structures either.

However, the fact that the implications and incompatibilities relating logically complex sentences in the vocabulary that logically extends particular material base vocabularies are substructural does not mean that the purely logical reason relations of those logically extended vocabularies—the implications and incompatibilities that hold in virtue of logic alone—must also be open-structured. We can extract or abstract such purely logical reason relations by an analogue of the Bolzano-Frege-Quine method of noting invariance under substitution referred to just above. Here what corresponds to observing which substitutions of nonlogical for nonlogical vocabulary preserve the goodness of sequents (what can be substituted *salva consequentia*) is to ask which sequents hold no matter what base vocabulary the logical connective rules of the sequent calculus NMMS are applied to. This will amount to only a tiny part of the sequents relating sets of sentences in the logically extended lexicon that hold relative to any particular material base vocabulary. But the reason relations that in this sense hold in virtue of the logical form of sentences, regardless of the content of the nonlogical sentences that appear there, will have the full topological closure structure that Gentzen and Tarski demanded of notions of logical consequence (and incompatibility as formal inconsistency). In other words, what we are offering is not a nonmonotonic (nontransitive, etc.) logic, as the recent tradition has conceived them. It is rather a logic for codifying nonmonotonic (nontransitive, etc.) consequence relations (and reason relations more generally). The expressivist philosophical reconceptualization of the issue—the story about what kind of logic is wanted once one acknowledges that reason relations do not in general have the strong closure structure of traditional purely logical reason relations—accordingly can be seen to have substantial consequences for the logics one crafts in response.

The striking fact is that in spite of it remaining tractable and retaining its substantial expressive power when applied to radically substructural base vocabularies, NMMS is in a clear sense, just classical logic. In the context of full global topological structural rules (Gentzen's Weakening and Cut, as well as Contraction and Contaiment), the mixed additive-multiplicative, reversible Ketonen rules of NMMS just yield classical logic. There are many such formulations of classical logic that are equivalent (in the sense of determining the same purely logical consequence relation) in the presence of those strong topological closure-structural restrictions, but that come apart if those restrictions are relaxed. It turns out that under such open-structured circumstances, the expressive powers of the various logics varies substantially. We have chosen carefully among those once-equivalent, now-diverging specifications of classical logic, to find one that best satisfies expressivist requirements.

3.5 Appendix

It will prove useful to start by introducing the notion of a proof-search.

Definition 20 (Proof-search). A (root-first) proof-search produces a proof-tree from a sequent $\Theta \succ \Sigma$, which is the root of the tree, by recursively applying the following procedure until the process terminates when the proof-tree no longer changes: (i) If $\Gamma \succ \Delta$ is the leaf of a branch of the tree at the current stage and all the sentences in Γ and Δ are atomic, then the branch remains unchanged. (ii) Otherwise, we look for the first complex sentence in $\Gamma \succ \Delta$ (starting on the left, ordering the sentences in Γ and Δ alphabetically) and build the branch up from that leaf by using the appropriate rule of NMMS. For example, we apply the top-to-bottom version of [L$\vee$] (moving upwards in the tree) if the left-most complex sentence in our sequent is a disjunction, thus yielding a fork in the tree with three new leaves that are the three top sequents of our application of [L$\vee$], and similarly for the other cases. (Although we work with sets (and so contraction is built in), we represent the sets in our sequents with the number of copies of sentences that we get by applying this procedure to the given representation of the root, thus treating our sets (in how we represent them) like multi-sets.)

The results of proof-searches for representations of the root sequent that differ in the numbers of copies of sentences do not differ, up to differences in the number of copies of sentences.

Proposition 21. *Proof-searches on* $\Gamma, A \succ \Delta$ *and* $\Gamma, A, A \succ \Delta$ *yield the same results, and the same holds for proof-searches on* $\Gamma \succ A, \Delta$ *and* $\Gamma \succ A, A, \Delta$, *up to differences in the number of sentences in the representations of the sets.*

Proof. If A is atomic, then the proof-search leaves it untouched. If A is complex, it is a conjunction, disjunction, or negation (treating the conditional as defined for simplicity, with $A \to B$ being $\neg A \vee B$). Suppose $A = B \wedge C$. Then applying our procedure with [L$\wedge$] to $\Gamma, B \wedge C \succ \Delta$ yields $\Gamma, B, C \succ \Delta$, and applying it twice to $\Gamma, B \wedge C, B \wedge C \succ \Delta$ yields $\Gamma, B, C, B, C \succ \Delta$. Thus, the resulting set of premises is identical, up to the number of copies of sentences. The cases for [R$\vee$], [L$\neg$], and [R$\neg$] are analogous. For [R$\wedge$], applying our procedure to $\Gamma \succ B \wedge C, \Delta$ yields $\Gamma \succ B, \Delta$ and $\Gamma \succ C, \Delta$ and $\Gamma \succ B, C, \Delta$. Applying the procedure twice to $\Gamma \succ B \wedge C, B \wedge C, \Delta$ yields $\Gamma \succ B, B, \Delta$ and $\Gamma \succ C, B, \Delta$ and $\Gamma \succ B, C, B, \Delta$ and $\Gamma \succ B, C, \Delta$ and $\Gamma \succ C, C, \Delta$ and $\Gamma \succ B, C, C, \Delta$ and $\Gamma \succ B, B, C, \Delta$ and $\Gamma \succ C, B, C, \Delta$ and $\Gamma \succ B, C, B, C, \Delta$. Each of the conclusion sets in these sequents is identical to that of one

of the three sequents above, up to differences in the number of copies of sentences. The case for [L$\vee$] is analogous. ■

Proposition 22. *Proof-searches (for finite sequents) terminate, and their results are the same if we change the order of the sentences in* Γ *and* Δ.

Proof. Proof-searches terminate because the root contains finitely many logical connectives, and the children of a node always contain one fewer connective than the parent node.

To show that the order doesn't matter, it suffices to show that for each pair of rules, the order in which they are applied doesn't matter. If we have, for example, $\neg A, B \vee C, \Gamma \succ \Delta$, applying our procedure to the first two sentences yields: $B, \Gamma \succ \Delta, A$ and $C, \Gamma \succ \Delta, A$ and $B, C, \Gamma \succ \Delta, A$. This result is the same whether we use [L$\neg$] first and then [L$\vee$] or the other way around. The same holds for all pairs of rules. Hence, the result of a proof-search is order-independent. ■

Proposition 23. *If all the leaves of a proof tree of* NMMS *can be weakened with any set of atomic sentences, then the root of the proof tree can be weakened with any set of sentences.*

Proof. By induction on the maximal complexity of the sentences with which one weakens the sequent. ■

Proposition 24. *If base* $\mathfrak{B}$ *obeys Containment, then for all* $\Gamma \cap \Delta \neq \emptyset$ *the sequent* $\Gamma \succ \Delta$ *is derivable in* $\mathsf{NMMS}_{\mathfrak{B}}$*, that is, the logical extension of the base also obeys Containment.*

Proof. By induction on the complexity of the most complex sentences in $\Gamma \cup \Delta$. The base case is immediate from the fact that $\underset{\mathfrak{B}}{\succsim}$ obeys Containment. For the induction step it suffices to note that all instances of Containment that feature sentences of complexity of at most $n+1$ can be derived from instances of Containment that feature sentences whose complexity is at most n. ■

Proposition 25. *If base* $\mathfrak{B}$ *obeys Containment and* $\Gamma \underset{CL}{\vdash} \Delta$ *holds in classical propositional logic, where the sentences in* Γ *and* Δ *are in the logical extension of the lexicon of* $\mathfrak{B}$*, then the sequent* $\Gamma \succ \Delta$ *is derivable in* $\mathsf{NMMS}_{\mathfrak{B}}$*, that is,* $\underset{CL}{\vdash} \subseteq \succsim$.

Proof. If [Weakening] and [Mixed-Cut] hold, then the rules of NMMS are equivalent to (the propositional fragment of) Gentzen's familiar LK sequent rules for classical logic. In Gentzen's LK, it is well-known that [Weakening]

can be absorbed into the axioms and the axioms can be restricted to atomic sequents, and [Mixed-Cut] can be eliminated by pushing applications of it up the proof tree to the leaves, where [Mixed-Cut] holds because the instances of Containment are closed under [Mixed-Cut]. ■

Proposition 26. *If* $\Gamma \cup \Delta \subseteq \mathfrak{L}_{\mathfrak{B}}$, *then* $\Gamma \mathrel{|\!\sim} \Delta$ *just in case* $\Gamma \mathrel{|\!\sim_{\mathfrak{B}}} \Delta$.

Proof. It suffices to note that all the rules of NMMS introduce logically complex sentences into the bottom sequent. ■

Proposition 27. *All the rules of* NMMS *are invertible, that is, the bottom sequent is derivable just in case all the top sequents are derivable.*

Proof. By induction on proof height. We just do the proof for [L→]; the other cases are analogous. Suppose our proposition holds for proofs of height n, and we derive $\Gamma, A \rightarrow B \succ \Delta$ is a proof of height $n+1$. If the root comes by [L→], we are done. If the root comes by any other rule, then $A \rightarrow B$ occurs on the left in all top sequents, which are derived in proofs of height n. Hence, we can apply our induction hypothesis, which yields three sequents for each top sequent of the last rule application of our original proof tree. And by applying the last rule applied in the original proof tree we can now derive the desired three sequents: $\Gamma \succ \Delta, A$ and $\Gamma, B \succ \Delta$ and $\Gamma, B \succ \Delta, A$. ■

Proposition 28. *If* $\Theta \cup \Lambda \subseteq \mathfrak{L}$ *and the sets are finite, then* $\mathfrak{L}$ *includes sentences that make explicit by implication that* Θ *is a reason for* Λ *and that* Θ *is a reason against* Λ*; i.e., there are sentences,* ϕ *and* ψ*, such that, for all* Γ *and* Δ*, we have* $\Gamma \mathrel{|\!\sim} \phi, \Delta$ *just in case* $\Gamma, \Theta \mathrel{|\!\sim} \Lambda, \Delta$*, and we have* $\Gamma \mathrel{|\!\sim} \psi, \Delta$ *just in case* $\Gamma, \Theta, \Lambda \mathrel{|\!\sim} \Delta$.

Proof. $\bigwedge \Theta \rightarrow \bigvee \Lambda$ meets the condition for ϕ. And $\neg \bigwedge \Theta \cup \Lambda$ meets the condition for ψ. For, the invertibility of the rules of NMMS ensures that $\Gamma \mathrel{|\!\sim} \bigwedge \Theta \rightarrow \bigvee \Lambda, \Delta$ just in case $\Gamma, \Theta \mathrel{|\!\sim} \Lambda, \Delta$; and $\Gamma, \Theta, \Lambda \mathrel{|\!\sim} \Delta$ if and only if $\Gamma \mathrel{|\!\sim} \neg \bigwedge \Theta \cup \Lambda, \Delta$. ■

Proposition 29. *For any sequent* $\Gamma \succ \Delta$*, with* $\Gamma \cup \Delta \subseteq \mathfrak{L}$*, there is a unique set* `AtomicImp` *of base-vocabulary sequents such that* $\Gamma \succ \Delta$ *is derivable in* $\mathsf{NMMS}_{\mathfrak{B}}$ *just in case* `AtomicImp` $\subseteq \mathrel{|\!\sim_{\mathfrak{B}}}$.

Proof. All the rules of NMMS are invertible. Hence, supposing that $\Gamma \succ \Delta$ contains finitely many sentences of finite length, the process of applying the NMMS rules repeatedly in a proof search on $\Gamma \succ \Delta$ yields a set of atomic sequents in finitely many steps. Since the order in which we apply these

(inverted) rules does not matter, this result is unique. Call the set of these atomic sequents $\texttt{AtomicImp}$. Then $\Gamma \succ \Delta$ is derivable in $\mathsf{NMMS}_{\mathfrak{B}}$ just in case $\texttt{AtomicImp} \subseteq \mathrel{\mid\!\sim}_{\mathfrak{B}}$. ■

Proposition 30. *$X, \Gamma \succ \Delta, Y$ is derivable for all $\langle X, Y\rangle \in \mathcal{P}(\mathfrak{L}_{\mathfrak{B}}) \times \mathcal{P}(\mathfrak{L}_{\mathfrak{B}})$ if and only if $Z, \Gamma \succ \Delta, U$ is derivable for all $\langle Z, U\rangle \in \mathcal{P}(\mathfrak{L}) \times \mathcal{P}(\mathfrak{L})$.*

Proof. The right to left direction is immediate from $\mathfrak{L}_{\mathfrak{B}} \subset \mathfrak{L}$. The left to right direction follows from Proposition 23. ■

Proposition 31. *If for every top sequent $\Gamma \succ \Delta$ of an application of a rule of* NMMS *(or* $\mathsf{NMMS}^{\backslash ctr}$*) $\Gamma \mathrel{\mid\!\sim}^{\uparrow} \Delta$ holds, then, for the bottom sequent $\Theta \succ \Sigma$, we have $\Theta \mathrel{\mid\!\sim}^{\uparrow} \Sigma$.*

Proof. We can weaken each top sequent with the sentences (on the left and the right) with which we want to weaken the bottom sequent, and derive the desired sequent using the original rule. ■

Proposition 32. *$\Gamma \succ^{\uparrow} \Delta$ is derivable in* $\mathsf{NMMS}^{\uparrow}$ *if and only $\forall \langle X, Y\rangle \in \mathcal{P}(\mathfrak{L}) \times \mathcal{P}(\mathfrak{L})$ $(X, \Gamma \mathrel{\mid\!\sim} \Delta, Y)$.*

Proof. Left-to-right: Suppose that $\Gamma \succ^{\uparrow} \Delta$. Then there is a $\mathsf{NMMS}^{\uparrow}$ proof tree in which all leaves are of the form $\Theta \succ^{\uparrow} \Lambda$. So we can weaken all leaves with arbitrary atomic sentences. By Proposition 23, it follows that $\forall \langle X, Y\rangle \in \mathcal{P}(\mathfrak{L}) \times \mathcal{P}(\mathfrak{L})$ $(\Gamma \mathrel{\mid\!\sim} \Delta)$.

Right-to-left: Suppose that $\forall \langle X, Y\rangle \in \mathcal{P}(\mathfrak{L}) \times \mathcal{P}(\mathfrak{L})$ $(X, \Gamma \mathrel{\mid\!\sim} \Delta, Y)$. Hence, $\forall \langle X, Y\rangle \in \mathcal{P}(\mathfrak{L}) \times \mathcal{P}(\mathfrak{L})$ $(X, \Gamma \succ \Delta, Y)$ is derivable in $\mathsf{NMMS}^{\uparrow}$. In particular, $\forall \langle X, Y\rangle \in \mathcal{P}(\mathfrak{L}_{\mathfrak{B}}) \times \mathcal{P}(\mathfrak{L}_{\mathfrak{B}})$ $(X, \Gamma \succ \Delta, Y)$ is derivable in $\mathsf{NMMS}^{\uparrow}$. In this case, X and Y cannot enter the proof tree anywhere but in the leaves. Since $\langle \emptyset, \emptyset\rangle \in \mathcal{P}(\mathfrak{L}_{\mathfrak{B}}) \times \mathcal{P}(\mathfrak{L}_{\mathfrak{B}})$, we know that $\Gamma \succ \Delta$ is derivable. It follows that, for each leaf of the proof tree for the root $\Gamma \succ \Delta$, we have the leaf $\Theta \succ \Lambda$ itself and also, $\forall \langle X, Y\rangle \in \mathcal{P}(\mathfrak{L}_{\mathfrak{B}}) \times \mathcal{P}(\mathfrak{L}_{\mathfrak{B}})$, the atomic sequent $X, \Theta \succ \Lambda, Y$. Hence, we have $\Theta \succ^{\uparrow} \Lambda$ for each leaf. So, we can simply decorate the proof tree for $\Gamma \succ \Delta$ with an upward arrow on which sequent, and this yields a proof tree for $\Gamma \succ^{\uparrow} \Delta$. ■

Proposition 33. *$\Gamma \succ \Box A, \Delta$ is derivable in* $\mathsf{NMMS}^{\Box}$ *if and only if $\Gamma \mathrel{\mid\!\sim}^{\uparrow} A, \Delta$.*

Proof. The right to left direction is ensured directly by our rules. $\Gamma \mathrel{\mid\!\sim}^{\uparrow} A, \Delta$ holds if $\Gamma \succ {}^{\uparrow}A, \Delta$ is derivable. And we can apply [R$\Box$] to get $\Gamma \succ \Box A, \Delta$.

For the left to right direction it suffices to show that [R$\Box$] is invertible and Proposition 32 still holds for $\mathsf{NMMS}^{\Box}$. The argument from the proof

of Proposition 27 can be adjusted to apply to [R□]. And the argument for Proposition 32 works exactly as for $\mathsf{NMMS}^{\uparrow}$. ■

Proposition 34. *If all the sentences in* $\Gamma \cup \Delta \cup \{A\}$ *are in the language of classical propositional logic, then the sequent* $\Gamma \succ (\!|cl|\!)A, \Delta$ *is derivable in* $\mathsf{NMMS}^{(\!|cl|\!)}$ *if and only if* $\Gamma \vdash_{CL} A, \Delta$.

Proof. Suppose that all the sentences in $\Gamma \cup \Delta \cup \{A\}$ are in the language of classical propositional logic. Then $\Gamma \vdash_{CL} A, \Delta$ holds just in case there is a derivation of $\Gamma \succ A, \Delta$ in $\mathsf{NMMS}^{(\!|cl|\!)}$ whose leaves are all instances of Containment. This happens just in case there is a derivation of $\Gamma \succ^{cl} A, \Delta$ in $\mathsf{NMMS}^{(\!|cl|\!)}$; and the rules for $(\!|cl|\!)$ cannot be used in such a derivation because $(\!|cl|\!)$ does not occur in the root sequent. And since [R$(\!|cl|\!)$] is invertible (by the same argument as the other rules), this happens just in case there is a derivation of $\Gamma \succ (\!|cl|\!)A, \Delta$. ■

Notes

1 While a discussion of alternative approaches to nonmonotonic consequence would lead us too far afield, we can illustrate our point in the main text with two quick remarks: First, in standard preferential logics, one adds a partial order over classical models and says that $\Gamma \mid\!\sim A$ if all the models of Γ that are minimal in the partial order are models of A (Kraus et al., 1990). Since there is no model of $\bot$, it follows that $\Gamma \mid\!\sim \bot$ if and only if there is no classical model of Γ, that is, if $\Gamma \vdash_{CL} \bot$. Thus, standard preferential logics must treat incoherence as classical and indefeasible.

Second, in standard default logics, one adds default rules of the form $\phi : \mathsf{M}\psi / \chi$ to classical logic, which are to be read as saying that given ϕ one can infer χ, as long as it is consistent to assume that ψ (Reiter, 1980). Then extensions of sets of sentences are defined by a procedure that adds the conclusion of a default whose premises are available, when this can be done in a way that yields a classically consistent result, closing the result under classical consequence. Such procedures can vary in their details (for instance, in the ordering which defaults are applied, etc.). And the consequences of a set of sentences are then defined in terms of such extensions; the simplest options are to take either the union (credulous strategy) or the intersection (skeptical strategy) of all such extensions. A classically consistent set of sentences cannot yield a classically inconsistent extension, as classical consistency is the chief constraint on extensions. Hence, the intersection of such extensions is also classically consistent. So, given a skeptical strategy, $\Gamma \mid\!\sim \bot$ if and only if $\Gamma \vdash_{CL} \bot$, as in preferential logics. The union of extensions can be classically inconsistent, and this might change if we add extra sentences to the set of sentences with which we begin. So, the credulous strategy can allow for the nonmonotonicity of incoherence. However, it is still classical consistency that is driving and

underlying this account, as the inconsistency of the conclusions arises only at the stage where we take the union of the classically consistent extensions.

2 Cross (2003) gives one of the rare discussions of nonmonotonic incompatibility. Let us point out that we avoid Cross's result that $\Gamma, A \mathrel{|\!\sim}$ if and only if $\Gamma \mathrel{|\!\sim} A$ yields a collapse into monotonicity because we reject the principle that if $\Gamma \mathrel{|\!\sim} A$ and $\Gamma \mathrel{|\!\sim} \neg A$, then $\Gamma \mathrel{|\!\sim}$. Given our negation rules below, this principle is equivalent to Cut and so we reject it when we reject transitivity. It is worth noting, however, that we accept the related principle that $\Gamma, A, \neg A \mathrel{|\!\sim} \Delta$, for any Γ, Δ, and A. We thus hold that explicit contradictions are persistently incoherent and imply everything. But if the contradiction is only implicit (in that both A and not-A are implied), then a distinction can still be made between what follows from the premises and what does not.

3 Below they will sometimes be multi-sets, namely when we discuss potential failures of contraction. Multi-sets are like sets except that multi-sets can agree in all their elements and still be distinct because they include different numbers of occurrences (numbers of copies) of those elements, for example, $\{A, A\}$ and $\{A\}$ are the same set but distinct multi-sets. We will flag when such complications become relevant.

4 These are of course merely the best known generalizations, there are also hypersequent calculi and many other generalizations. The generalization to multi-sets is the only one that will become relevant below.

5 We will write interchangeably of accepting sentences and accepting their contents. We will not offer any treatment of expressions that are sensitive to contexts of utterance or the like; so that we can make the simplifying assumption that sentences always express the same content.

6 There are also formulations of bilateralism in the literature that do not work with multiple conclusions (see Rumfitt, 2000; Smiley, 1996). The conception of consequence in these other formulations differs from ours. So we set them aside, although a detailed comparison would be an interesting project.

7 We think of this interpretation of offering a natural connection between our use of sentences and consequence relations. We thus disagree with Steinberger's (2011) claim that semantic inferentialism is in tension with multiple-conclusion calculi.

8 We will sometimes assume that the sets among which consequence relations hold are finite. However, we make this assumption for convenience. We are open to generalizing our notions of consequence, sequents, proof trees (below), and the like to infinite sets of sentences and "proof trees" of infinite height (and breadth). Since taking this complication into account would lead to technical complications that we wish to avoid here (such as the need for transfinite inductions), we make the simplifying assumptions that implications hold among finite sets and that proof trees are finite in height. Of course, this assumption would have substantive implications if we wanted to add arithmetic to our consequence relations. However, we will not consider arithmetic or other theories for which this assumption will be crucial.

9 This approach was first developed in (Hlobil, 2016), but there it was applied in a Set-Formula setting to yield super-intuitionistic logics. Our use here is closer

to (Kaplan, 2018), where the idea is applied in a Set-Set setting to yield super-classical logics. Applications that yield first-order logics and relevance logics have also been developed (Shimamura, 2019, 2017; Hlobil, 2018).

10 Note that the conditional of NMMS is as much an idealization and simplification of the natural language conditional as the so-called "material" conditional. Like the "material" conditional, our $\phi \to \psi$ is equivalent to $\neg\phi \vee \psi$. Hence, denying the conditional "It is raining → The sky is blue" is out-of-bounds if and merely because it is out-of-bounds to assert "It is raining" while denying "The sky is blue." And, by an inverted application (to be discussed shortly) of [L→], if it is out-of-bounds to assert "It is raining → The sky is blue," then it is also out-of-bounds to deny "It is raining" or to assert "The sky is blue" or to do both. Our aim regarding the conditional here is not to offer a plausible formal codification of the natural language conditional but rather to introduce a conditional that can make explicit reason relations.

11 It can easily be shown that negation together with one of the other connectives suffices to define the remaining two connectives, in a way that is broadly analogous to the functional completeness of the connectives in classical logic.

12 We will not speculate about how this might work at a cognitive or pragmatic level, which would concern only particular agents and particular communities. What matters to us is merely that nothing more than an appropriate ability to use the base language is needed, whatever the mechanism is to transform this ability into the ability to use logical vocabulary.

13 In standard technical language: NMMS has the subformula property.

14 Notable exceptions are connexive logics and strict-tolerant logic. We return to the latter below.

15 For our purposes in this book, we can define the complexity of a sequent (or a sentence) as the number of connectives that occur in it. Sometimes a more careful definition as the number of connectives embedded under which the most deeply embedded atomic sentence occurs are more convenient. However, what we say holds for all such definitions of complexity. Hence, we will often not be precise about which notion of complexity we have in mind.

16 We omit quotation marks to avoid clutter. Strictly speaking, what flanks the snake-turnstile are terms for object language sentences and not object language sentences. We trust context to disambiguate between a sentence and its quote-name.

17 Here we use a recipe for turning examples of monotonicity failures into examples of transitivity failures that was developed, in the context of discussing conditionals, by Ryan Simonelli (2022). Related observations can be found in Nair (2019).

18 Lakatos (1976) might be read as giving an example of how a definition of a polyhedron that avoids such failures of monotonicity and transitivity could arise. The simplistic example regarding triangles is alluded to in Lakatos (1976, 24), when student Sigma suggests that there are some mathematical theorems with exceptions and others without exception, such as "the angle sum of all plane triangles is always equal to two right angles." And the passage continues: "Epsilon (to Kappa): Who is this muddlehead? He should learn something about

logic. Kappa (to Epsilon): And about non-Euclidean plane triangles." It seems to us that the temptation to make such objections shows how powerful the song of the siren of Logicism about reasons is. Our project is to sail past such sirens, tied to the mast if necessary, to explore what lies on the other side.

19 This can happen because, in nontransitive consequence relations, mutual implication of two sentences does not entail that the two sentences can be replaced for each other as premises and conclusions *salva consequentia*. The importance of this fact will become clearer in Chapter Five, when we introduce the notion of conceptual content in terms of substitutability *salva consequentia*.

20 The restriction to finite sets of base sequents holds because it might take infinitely many steps to "combine" infinitely many base sequents into one logically complex sequent. If we are willing to allow for proof trees with infinitely many steps in our sequent calculus, then the theorem holds for all sets of implications in the base consequence relation.

21 A validity predicate that strongly represents the reason-for relation in the validity-free fragment of the language is presented by Hlobil (2020).

22 As we will see in later chapters, this idea applies in important ways to "candidate implications" or "bad implications"; their range of subjunctive robustness is the set of pairs whose addition turns these implications into good implications. This idea will become important in the implication-space semantics below. The definition we give here is already general in this way, although we do not stress this here.

23 It is possible to incorporate into the sequent calculus set theoretic manipulations of the set in the place here occupied by R and Z. This is important in the case of super-intuitionistic logics (see Hlobil, 2016). To quickly illustrate the basic idea: If we define the upward arrow by saying that $\Gamma \mid\!\sim^{\uparrow Z} \Delta$ if and only if $\forall \langle \Theta, \Sigma \rangle \in Z\ (\Theta, \Gamma \mid\!\sim \Sigma, \Delta)$, then the following are immediate consequences: If $\Gamma \mid\!\sim^{\uparrow U} \Delta$ and $\Gamma \mid\!\sim^{\uparrow W} \Delta$, then $\Gamma \mid\!\sim^{\uparrow U \cup W} \Delta$. And if $U \subseteq \mathtt{RSR}$ of $\langle \Gamma, \Delta \rangle$, then $\Gamma \mid\!\sim^{\uparrow U} \Delta$ and $\Gamma \mid\!\sim^{\uparrow U \cap V} \Delta$ for any set V. Such consequences can then inform new kinds of sequent rules (see Hlobil, 2016). We will not need this machinery in this work. Hence, we will not introduce any details here.

24 Everything in this section applies to both calculi. To avoid clutter, we will henceforth refrain from mentioning $\mathsf{NMMS}^{\backslash ctr}$ explicitly.

25 Here and elsewhere we set aside the structural principle of Permutation. We consider Contraction only occasionally. Thus, we usually take for granted that collections of premises and conclusions are, in effect, sets.

26 That is not usually done because it blocks showing admissibility of Cut, given failures of Monotonicity. But we do not want that, since not all the base vocabularies whose implication relations we want to make explicit satisfy the strong (mixed-context) transitivity condition that is Cut.

4 Truth-Taking and Truth-Making
How They Share a Rational Form

Ulf Hlobil

So far, we have pursued a pragmatics-first approach to reason relations. We began with the practice of giving reasons for and against claims. This practice is constrained by reason relations between the contents of claims; and we often let relations among sentences with those contents go proxy for the relations among contents.[1] In the first instance, we understood such reason relations to hold among sentences that do not contain logical vocabulary, that is, among atomic sentences. And we understood logic to be concerned with the norms that govern discourse in two ways: First, we introduced logical vocabulary by specifying the role of the logical lexicon in reason relations, namely by giving sequent rules. Second, we claimed that it is the characteristic job of logical vocabulary to make explicit reason relations, and hence to make explicit constraints on discursive norms governing the use of sentences.

In this chapter, we turn to a semantics-first approach to reason relations. Following such an approach, one aims to give an account of reason relations in terms of relations among what is represented by sentences. Hence, a semantics-first approach to reason relations must start with an account of what and how sentences represent. Our aim in this chapter is to show that there is a way to think about the sentences of our language as representing how things are, in such a way that the account is isomorphic to the pragmatics-first approach presented in the previous chapters.[2] This isomorphism is at the heart of the account of representation that we develop below. And the relation between metavocabularies for reason relations that the isomorphism specifies is at the heart of this book.

The upshot of this chapter is that we can give an account of the reason relations among sentences either in a pragmatic-normative metavocabulary, by specifying the norms that govern the use of the sentences, or we can give an account of these reason relations in a semantic-representationalist metavocabulary, by specifying the modal relations among states represented by sentences. Crucially, however, we can develop these two

accounts in such a way that they are isomorphic. We thus arrive at an account of how assertions and denials of sentences—or the acceptance and rejection of claims—represent worldly states that is isomorphic to the account of the norms that govern the uses of sentences. And, as we shall see, this underwrites a version of conceptual realism.

4.1 Rational Forms

In this section, we put in place a semantics-first approach to reason relations at a conceptual level. We codify this approach formally in the truth-maker theory discussed in the following sections.

4.1.1 Representation and Truth

A semantics-first approach to content and reason relations starts with the idea of representation, which is closely related to the idea of truth. What is true or false are contentful items, such as (uses of) sentences, or perhaps mental states or acts.[3] For the sake of simplicity, we will restrict ourselves to sentences as contentful items and, hence, as vehicles of representation. But what we say is meant to apply to other vehicles of representation as well.[4] For our current purposes, a vehicle of representation is anything that can be true or false. So, in keeping with our focus on the sentential level, the kind of representation that is our topic here is not the representation of objects or properties but rather the representation of whatever it is that declarative sentences represent.

What does it mean for a sentence to be true or false? We can understand the truth and falsity of a sentence in terms of the accuracy of uses of the sentence in assertions and denials. To assert that things are thus-and-so while things are thus-and-so or to deny that things are thus-and-so while things are not thus-and-so is accurate. To assert that things are thus-and-so while things are not thus-and-so or to deny that things are thus-and-so while things are thus-and-so is inaccurate. A sentence is true if its assertion is accurate and its denial is inaccurate, and a sentence is false if its assertion is inaccurate and its denial accurate.[5]

A sentence represents things to be a particular way if how these things really stand serves as a standard of truth and falsity for the sentence. Thus, there is a norm governing representations that deems them true or false according to how things stand with what is represented. We may call this essential feature of representation "Normative Governance."

Normative Governance
If ϕ represents x, then whether ϕ is true or false depends on how things stand regarding x.

It follows from Normative Governance that what a sentence represents is, at least in part, a matter of the conditions under which the sentence is true and the conditions under which it is false. There are, however, different ways to spell out this idea of conditions of truth and falsity. Since we are interested in open reason relations and, hence, in frameworks that allow us to reject global structural closure principles, we must be careful to think about conditions of truth and falsity in a fine-grained and flexible way. It will become clear below that we can do so in the framework of truth-maker theory. In truth-maker theory, one thinks about the conditions of truth and falsity of a sentence in terms of states that make a sentence true or false.

Truth-maker theory does not commit one to any particular account of the states that are truth-makers and falsity-makers. They could be events, states of affairs, objects, properties, and more. It is even often convenient to have models in which states are (sets of) sentences. As we shall see below, all that is formally required of states is that they are partially ordered, that every subset of them has a least upper bound with respect to this partial order, and that there is a partition among them into the possible states and the impossible states.

In a philosophical account of representation, however, we should say something about what we take states to be, beyond the very thin requirements of the formalism. To understand what states are, note that a false sentence represents things as being a certain way while that is not how things really are. This way for things to be, while they might actually be otherwise, is a kind of state. It is a kind of state and not a particular state, because there might be many specific ways in which things could be as the sentence represents them to be. The states of this kind are such that the sentence is true if such a state is a part of how things really are, that is, if a state of this kind obtains.[6] These states are the truth-makers of the sentence. Similarly, the states that are such that the sentence is false if they obtain are the falsity-makers of the sentence. To assert a sentence is to purport to state a fact, because it is to convey that one of the states that would make the sentence true is part of how things really are. And to deny a sentence is to purport to state a fact, because it is to convey that one of the states that make the sentence false is part of how things really are.

It is a tricky question what the falsity-makers of many sentences are. What is, for instance, a falsity-maker of "Wittgenstein was French"? Is it perhaps the state of Wittgenstein having been Austrian? Or do we have to posit a negative state of Wittgenstein failing to be French? Or is it the state of reality failing to contain a truth-maker for "Wittgenstein was French"? We will leave such questions unanswered, as all that matters for our purposes is that a falsity-maker is whatever makes the denial of a sentence accurate. It is worth noting, however, that many states leave the truth or falsity of many sentences undecided. Plausibly, the state of Socrates

having been Greek, for instance, makes the sentence "Wittgenstein was French" neither true nor false. In fact, our notion of a state is such that the state of Socrates having been Greek does not include anything that is not relevant for whether Socrates was Greek. In Kit Fine's terminology, our truth-makers and falsity-makers are *exact* truth-makers and *exact* falsity-makers. It is also worth noting that we do not exclude impossible states, such as the state of two plus two being five. These are respects in which our notion of conditions of truth and falsity is fine-grained and differs, for instance, from the notion of possible worlds.

At this point, we should distinguish between two senses of "to represent." There is a sense of "to represent" according to which a true sentence represents the obtaining worldly state(s) in virtue of which the sentence is true, and a false sentence represents the obtaining worldly state(s) in virtue of which the sentence is false. Let us call this "factive representation." In this factive sense, only states that obtain can be represented; non-obtaining states cannot be represented. But we can also talk about what a sentence represents without presupposing that it is true or that it is false, thus allowing that the states that we say the sentence represents might not obtain. Let's call this "non-factive representation."[7] It will prove convenient to introduce the label "worldly proposition" for what a sentence represents in this non-factive sense. The worldly proposition that a sentence represents is given by the truth-makers and the falsity-makers of the sentence, and we can think of it, by convention, as the pair whose first member is the set of the sentence's truth-makers and whose second member is the set of its falsity-makers.[8] What a sentence factively represents is contained in what it non-factively represents (either as a truth- or as a falsity-maker) but not vice versa.

Let us sum up what we said so far. We said that a sentence (non-factively) represents the worldly proposition that is the pair of the sentence's truth-makers and the sentence's falsity-makers. This is a notion of representation, because it makes how things really are the standard for the truth or falsity of the sentence, thus underwriting Normative Governance. Moreover, the account connects with the pragmatics of using the sentence because an assertion of the sentence is accurate if and only if the asserted sentence is true, and a denial is accurate if and only if the asserted sentence is false. Note, however, that so far we have not said anything about what makes certain states a sentence's truth-makers and others its falsity-makers. In other words, we have so far not addressed the question: What makes a particular sentence represent a particular worldly proposition? Any representationalist, semantics-first account of reason relations and contents is crucially incomplete and underdeveloped unless and until it answers such metasemantic questions.

The question of what makes some states truth-makers and other falsity-makers of a sentence can be understood in two different ways. According to the first way of understanding the question, it asks what it takes for a sentence to represent a particular worldly proposition: What counts as a sentence representing a worldly proposition? According to the second way of understanding the question, it asks how a sentence acquires and keeps a worldly proposition as that which it represents: How do speakers manage (what do they have to do) to make their sentences (continue to) represent particular worldly propositions?

We will be mainly concerned with the first of these two questions. Before addressing these questions, however, it is worth noting that we can change what a sentence represents by changing the norms that govern the correct use of the sentence. If we all started, for instance, to assess uses of the sentence "Fido is a cat" regarding its truth, justification, entailments, and the like (along with the further changes that this might require, for instance, in assessing uses of other sentences) as we now assess uses of the sentence "Fido is a dog," then the worldly proposition that is represented by "Fido is a cat" would change and would be the worldly proposition that is now (actually) represented by "Fido is a dog." Hence, what makes a particular sentence represent a particular worldly proposition must be a matter of the norms that govern the use of sentences. This is one place where the pragmatics-first approach from the previous chapters and the semantics-first approach that we are pursuing here are linked.

The question what it takes for a sentence to represent a particular worldly proposition thus turns into the following question: What must the norms that govern the use of sentences be like in order for a particular sentence to represent a particular worldly proposition?

The answer that we shall articulate in this chapter is that a sentence represents a worldly proposition just in case the sentence and the worldly proposition share what we call their "rational form."[9] A sentence and a worldly proposition share their rational form if a certain isomorphism holds between uses of the sentence and the states in a worldly proposition. More specifically, the relevant isomorphism must relate assertions of a sentence and its truth-makers as well as denials of a sentence and its falsity-makers, and the isomorphism holds with respect to the modal profile of assertions/denials and truth-/falsity-makers.

The question how a sentence acquires and keeps a worldly proposition as that which it represents, correspondingly, turns into the question: How do the norms that govern the use of sentences emerge and develop so as to make particular sentences represent particular worldly propositions?

The general shape that an answer to this question should take, we think, is that the norms that govern the use of sentences must be sensitive to and

shaped by what these sentences represent. We call this idea "Covariant Tracking."

> *Covariant Tracking*
> If y represents x, then y would be different if x had been relevantly different.

The general idea behind Covariant Tracking is that in order for y to represent x, the representation y must tell us something about x. And it is not clear how this could be the case if y would be as it actually is independently of what is the case regarding x.[10] It makes a big difference, however, whether what we mean by "represents" in Covariant Tracking is factive representation or non-factive representation. If we mean factive representation, then the most plausible reading of Covariant Tracking is this: If sentence y factively represents its truth-maker (falsity-maker) x, then agents would not tend to assert (deny) y (to the extent that they do) if x had not obtained. Let us call this "Epistemic Covariant Tracking." If we understand Covariant Tracking as a condition on non-factive representation, however, the most plausible reading of it is very different, namely something like the following: If sentence y represents the worldly proposition x, then the norms governing assertions and denials of y would be different if something relevant about the worldly proposition x had been different. Let's call this "Semantic Covariant Tracking." As we just formulated it, it contains parameters that would need to be specified: What kind of norms governing assertions and denials are at issue? What are the relevant aspects of worldly propositions that would make a difference to these norms? As will become clear below, we hold that the answers to these questions bring in normative and alethic relations of incompatibility.

Epistemic and Semantic Covariant Tracking are related in complex ways. Fortunately, we can avoid these complications for our current project. For our purposes, it is Semantic Covariant Tracking that matters. And what Semantic Covariant Tracking requires is that if a sentence represents a worldly proposition, then the rational form of the sentence would be different if the rational form of the worldly proposition had been different. We will simply assume that this condition holds for the sentences that we are discussing. We, hence, assume that we have a way to adjust the norms that govern our sentences to the rational form of the worldly propositions that we aim to represent with our sentences. However, we leave a discussion of how we do this for another occasion.

To sum up, our view can be described as hylomorphic conceptual realism. The view is hylomorphic because it takes representation to be a relation between hylomorphic compounds that share their form but differ in their matter. The sentence that is representing and the worldly proposition that

it represents share their form, namely their rational form. The matter of a sentence consists in tokenings of what Sellars (1963) calls a "sign design" (as assertions or denials), while the matter of a worldly proposition is the matter of worldly states. The rational form is what makes tokenings of sign designs into an occurrence of a particular contentful sentence, and it is also what makes the matter of which a worldly state consists into an occurrence of a particular worldly proposition. The view is a version of conceptual realism because it says that conceptual structure occurs in reality as it is independently of our representation of it. For, to have a particular conceptual structure is to be informed by a particular rational form, and that rational form can be found not just in that which represents but also in that which is represented. When we accurately assert that things are thus-and-so, then what we say and what is the case is the same, in the sense that what we say and what is the case are occurrences of the same rational form, and *mutatis mutandis* for denials. Thus, the conceptual structure, the conceptual *morphe*, is realized in discursive acts and in worldly states.

The remainder of this chapter is an articulation of what it takes for a sentence and a worldly proposition to share a rational form. We will thereby provide a partial explanation of what rational forms are. The answer that we articulate here and defend over the course of the remaining chapters has two key components, namely functionalism and modalism about rational forms. Let us look at these two key components more closely.

4.1.2 Functionalism and Modalism

We maintain that, in virtue of having a particular rational form, that which has the rational form has a particular modal profile. Functionalism and modalism about rational forms concern this modal profile. In order to explain this, let us start with the traditional notion of a substantial form.[11]

The traditional notion of a substantial form is something that can inform different bits of matter, thereby making these bits of matter into unified wholes of a particular kind. Particular substances are what they are in virtue of having the substantial form that they have. Moreover, that something has a certain form has implications about what is possible and impossible for that thing. For example, if two bits of matter share the form of a strawberry, they are, in virtue of having that form, two strawberries, and certain things are possible and others impossible for strawberries. It is, for instance, impossible for a strawberry to be ripe but green, but it is possible for a strawberry to be larger than a grape.

Rational forms also have these traditional features of substantial forms: When a rational form informs some matter, this yields a unity of a specific kind with a particular profile of what is possible and impossible for it. If a rational form informs a tokening of a sign design, then it makes it into

an occurrence of a particular contentful sentence. And if a rational form informs a bit of the kind of matter that makes up worldly states, then it makes it into an occurrence of a particular worldly proposition.[12]

Functionalism about rational forms is the thesis that it is necessary and sufficient for something to have a particular rational form that it plays a particular role with respect to other things that have rational forms. Modalism about rational forms is the thesis that the role that something must play in order to have a particular rational form is a role in modal relations. These two theses together yield the claim that what it is to have a particular rational form is to have a particular modal profile with respect to other things that also have rational forms.

The two kinds of bearer of rational forms that are central to our discussion here are sentences and worldly propositions. For we want to know what it takes for a sentence and a worldly proposition to share their rational form. We hold that rational forms give their bearers, that is, the hylomorphic compounds, a very special modal profile. Indeed, to have a particular rational form just is to have a particular modal profile. And in order to understand this modal profile, we must note that bearers of rational forms occur in one of two opposing ways, which we may call "positive" and "negative" occurrences of the rational form. The positive occurrence of a rational form in a sentence is an assertion of the sentence, and a negative occurrence is a denial of the sentence. A positive occurrence of a rational form in a worldly proposition is a truth-maker of the worldly proposition, and a negative occurrence is a falsity-maker of that worldly proposition.

Positive and negative occurrences of rational forms stand in relations of compatibility and incompatibility.[13] What it means for two occurrences of rational forms to be incompatible is that their combination is impossible. The modal profile that bearers of rational forms have consists in the compatibility or incompatibility of their positive and negative occurrences with further occurrences of rational forms, in the same kind of matter. Thus, the modal profile of two bearers of rational forms is identical just in case their positive and negative occurrences are compatible and incompatible with exactly the same further occurrences of bearers of rational forms. That is, the modal profile of two bearers of rational forms are identical just in case the respective combinations of their positive and negative occurrences with other occurrences of rational forms are always either both possible or both impossible.

It will become important below that the kind of impossibility with respect to which a bearer of a rational form plays its role can differ with the matter in which the rational role occurs. In particular, when a rational form occurs in the tokening of a sentence, the kind of impossibility is normative. When the form occurs in a worldly state, however, the

relevant kind of impossibility is alethic. This suggests that what it takes for a sentence to have the rational form of a worldly proposition is that assertions of the sentence are normatively incompatible with exactly the assertions (denials) of those sentences whose truth-makers (falsity-makers) are alethically incompatible with the truth-makers of the worldly proposition, and denials of the sentence are normatively incompatible with exactly the assertions (denials) of those sentences whose truth-makers (falsity-makers) are alethically incompatible with the falsity-makers of the worldly proposition.

In order to formulate this idea more precisely, let upper case Latin letters range over sentences. Let $+(\phi)$ stand for an assertion of ϕ, and $-(\phi)$ for its denial. Moreover, let us write $|\phi| = \langle |\phi|^+, |\phi|^- \rangle$ for a worldly proposition that is mapped to sentence ϕ, where $|\phi|^+$ is the set of the sentence's truth-makers and $|\phi|^-$ is the set of its falsity-makers. Then we can formulate our thesis as follows.

Sharing Rational Forms

The function $|\cdot|$ maps sentences to worldly propositions with which they share their form if and only if, for every sentence S and every two collections of sentences $\{P_1, ..., P_l\}$ and $\{N_1, ..., N_k\}$,

(pos) any combination of $+(S)$ with $+(P_1), ..., +(P_l)$ and $-(N_1), ..., -(N_k)$ is normatively impossible if and only if any combination of states that contains one state from each of $|S|^+$ and $|P_1|^+ ... |P_l|^+$ and $|N_1|^- ... |N_k|^-$ is alethically impossible, and

(neg) any combination of $-(S)$ with $+(P_1), ..., +(P_l)$ and $-(N_1), ..., -(N_k)$ is normatively impossible if and only if any combination of states that contains one state from each of $|S|^-$ and $|P_1|^+ ... |P_l|^+$ and $|N_1|^- ... |N_k|^-$ is alethically impossible.

This can seem rather complex at first, but the idea is simple. For sentences to represent particular worldly propositions, there must be an isomorphism between collections of assertions and denials, on the one side, and collections of truth-makers and falsity-makers, on the other side. In particular, the isomorphism must map assertions of a sentence to truth-makers of the sentence, and it must map denials of the sentence to falsity-makers of the sentence. And this isomorphism must preserve the modal status of the collections it maps to one another. That is, normatively impossible combinations of assertions and denials are mapped to alethically impossible combinations of truth-makers and falsity-makers, and vice versa.

This account can explain why changing the norms that govern the use of a sentence can change the rational form of the sentence. By changing which assertions and denials are normatively incompatible with assertions and denials of a given sentence, we change the rational form of the sentence. Thus, we can make a sentence share the rational form of a worldly proposition and, so, make it represent that proposition by changing the norms that govern assertions and denials of the sentence in the appropriate way.

We can express the appeal to two kinds of modality by saying that our position is a *bimodal* hylomorphic conceptual realism.[14] By this we mean that the conceptual structure of rational forms can inform two kinds of matter that are distinguished by the modal relations that hold among the respective hylomorphic compounds. When rational forms inform matter that occurs in discursive acts, such as tokenings of sentences, then the modality in which occurrences of the forms stand is normative. And when rational forms inform matter that occurs in worldly propositions, such as worldly states, then the modality in which occurrences of forms stand are alethic. Thus, while modal roles define the forms of hylomorphic compounds, it is the matter that decides in what kind of modality these modal roles are realized.

To sum up, functionalism about rational forms says that rational forms are identical just in case their occurrences play the same role with respect to other occurrences of rational forms, in the same kind of matter. And modalism says that to play a particular role, in the sense at issue, is to have a particular modal profile. The two theses together say that what it is to have a particular rational form is to have positive and negative occurrences that are incompatible with certain other occurrences of rational forms. This view is functionalist because it says that what matters for whether something has a particular rational form is not its internal structure but merely the role it plays in a larger whole of similar items. That explains why rational forms can inform very different matter, namely the matter in which represent*ings* occur, such as tokenings of sign designs, and the matter in which represent*eds* occur, like worldly states.

4.1.3 Some Worries about Representation

We said that what it takes for a sentence to represent a particular worldly proposition is that they share their rational form. And we said that what rational form a sentence has must be a matter of the norms that govern its use. Now some noteworthy details of this account of representation start to emerge, and so we may consider some potential objections and replies as a way to further clarify our view.

A sentence represents a worldly proposition just in case they play the same modal role with respect to further sentences and worldly propositions, respectively, and this holds in virtue of appropriate Semantic Covariant Tracking. The sentence and the worldly proposition play the same modal role just in case assertions of the sentence and truth-makers of the worldly proposition as well as denials of the sentence and falsity-makers of the worldly proposition play, respectively, the same modal roles with respect to further assertions/denials and truth-/falsity-makers. This is the claim that we formulated in Sharing-Rational-Forms above.

A first potential problem for any account of representation is that it must make room for misrepresentation. This is easy for our account. Misrepresentation does not occur between sentences and the worldly proposition they represent but between positive and negative occurrences of the sentence and obtaining states in a worldly proposition. An assertion of the sentence misrepresents how things are just in case it is inaccurate, that is, if no truth-maker but rather a falsity-maker of the worldly proposition obtains. And a denial of the sentence misrepresents how things are just in case it is inaccurate, that is, if no falsity-maker but rather a truth-maker of the worldly proposition obtains.[15] By contrast, an assertion represents correctly how things are and, hence, states a fact just in case the assertion is accurate.

A second, and very old, worry about theories of representation according to which what represents must share something with what is represented is that the relation of sharing something is symmetrical, while the relation of representation is not symmetrical. Hence, we must say why a sentence represents a worldly proposition and not vice versa.

Our answer to this second worry is that the asymmetry between representing and represented can be explained in our account by the difference between normative and alethic modality. Occurrences of rational forms where the relevant kind of modality is normative are occurrences that represent, and occurrences of rational forms where the relevant kind of modality is alethic are occurrences that do not represent but can be represented. This is so because the difference between representing and represented is the difference between what is measured and the measure. What is represented is the measure against which the representing is measured: it is the standard for assessments of its correctness. And what is governed by a norm—by a measure—is what is measured. Hence, representings are governed by norms, and in particular the norms pertaining to representational correctness. In contrast to this, representables are not, merely in virtue of being representables, governed by such norms pertaining to representational correctness. The incompatibility relations among them are not normative. There are, of course, complicated cases, such as the representation of representings and alethic modal

relations among representings. Such complicated cases are outside of the scope of our foundational project in this book, and we leave their discussion for future occasions. As a response to the current worry, it suffices to say that when rational forms occur in a matter that requires a normative modality as the modality in which they have their characteristic modal profiles, then they inform representings. And when rational forms occur in a matter that requires an alethic modality for their modal profiles, then they inform what can be represented.

A third worry might be that there are too many isomorphisms to uniquely determine a representation relation. We may, for example, envisage a language with just two sentences such that their assertions are incompatible and the denial of either one is compatible with the assertion and with the denial of the other. One might think that these two sentences could represent any number of worldly propositions whose truth-makers are incompatible and whose falsity-makers are compatible, such as the worldly propositions represented in English by "Fido is a dog" and "Fido is a fish" or "This ball is red" and "This ball is green."

Our answer to this worry is that we must distinguish two issues. The first issue is a very general potential problem for any theory of representation. If a structure has some nontrivial automorphism, then any requirement that a certain structure-preserving mapping holds between this structure and another structure does not determine a unique mapping between the two structures. Hence, there cannot be any theory of representation that determines a unique mapping between representings and representeds, that characterizes a structure of representings in such a way that the structure has a nontrivial automorphism, and that characterizes the representation relation by saying that a certain structure-preserving mapping holds between the structure of representings and the structure of representeds.[16] There are three possible ways to address this issue: (a) One might hold that a theory of representation need not determine a unique mapping between representings and representeds. (b) One might try to characterize the structure of representings in a way that does not allow for nontrivial automorphisms. (c) One might include in one's theory of representation characterizations of representings that go beyond the structure of representings as it is in isolation from what is represented. This issue and its three possible approaches to it are not specific to our account of representation, and it seems to us that we are free to adopt any of the three approaches to it.

As philosophers who press the kind of worry that we are currently considering typically prefer an approach to the issue along the lines of (c), let us sketch how this approach could be integrated to our account of representation. Notice that whether an assertion of "This is a dog" is incompatible with an assertion of "This is a fish" depends on whether the

speaker points at the same thing while uttering the two tokens of "this." Hence, by our modal functionalism about rational forms, the rational form of "This is a dog" when asserted while pointing at one object is different from the rational form of "This is a dog" when asserted while pointing at another object. So, strictly speaking, we should not talk about the rational forms of sentences but rather about the rational forms of sentences in a particular context of use. In this way, which rational forms occur in what representings may depend on features of the representings that go beyond the structure of representings as it is in isolation from what is represented. We do not commit ourselves to this approach along the lines of (c) here. We merely wish to point out that the issue is a general one, and our account can be fleshed out and extended in ways that is analogous to familiar approaches to the issue.

The second issue is how we should think about potential but nonactual sentences of the language. To see what we mean, consider the example from above again, where we stipulated that our language contains just two sentences whose assertions are incompatible but no other occurrences of them in assertions and denials are incompatible. Are assertions or denials of these sentences incompatible with the assertions or denials of sentences that we might add to our language? If we answer that we meant our stipulation in such a way that the uses of these sentences are compatible with any uses of other sentences that we might add to our language, then what they represent is not uniquely determined only if there are at least two pairs of sets of states (worldly propositions) such that any combination of two states such that one comes from the first member of the first pair and the second from the first member of the second pair are incompatible and the states are otherwise compatible with all other states. The example of the worldly propositions expressed, in English, by "Fido is a dog" and "Fido is a fish" and again by "This ball is red" and "This ball is green" is not an example of such a case. Rather, the structure of alethic incompatibility (and not just the structure of normative incompatibility) has nontrivial automorphisms in such a case, and we will come back to this momentarily. If, however, we answer that we meant our stipulation in such a way that the uses of our two sentences may or may not be compatible with the uses of other sentences that we might add to our language, then the question arises whether such an extension of our language must preserve the isomorphism between normative and alethic incompatibility. We hold that the isomorphism must indeed be preserved by all such possible extensions in order for a representation relation to hold. That is how our account of representation ought to be understood. Hence, in order to determine the isomorphism that explains representation, we must look at all possible extensions of a given language. If two distinct isomorphisms between normative and alethic incompatibilities remain even when we take arbitrary

extensions of our language into account, then there must be a nontrivial automorphism in the structure of alethic incompatibilities. Therefore, such an automorphism in the structure of alethic incompatibilities must exist for either one of our two answers to yield a potentially problematic underdetermination of the representation relation. Such an automorphism would be an extremely interesting symmetry in the modal structure of reality. If there is such a symmetry, then there is a corresponding symmetry in the normative incompatibilities. And this automorphism in the normative incompatibilities is responsible for the fact that our account cannot pick out a unique worldly proposition for each sentence as what it (non-factively) represents. In such a case, we are open to the possibility that there is no fact of the matter which of the two isomorphisms constitutes a representation relation. It is difficult to find any realistic example of such an automorphism. Perhaps the automorphism between the mathematics of complex numbers and that same structure with i and $-i$ replaced for each other is an example. If so, we would be open to the possibility that it is indeterminate whether our sentence "$i^2 = -1$" represents the worldly proposition whose truth-maker is that the square of i is -1 or the worldly proposition whose truth-maker is that the square of $-i$ is -1. This seems to us to be a kind of underdetermination of the representation relation that may be acceptable, should it really turn out to be intelligible.

As a final remark regarding the third worry, we want to point out that there may be cases of genuine indeterminacy of which worldly proposition is nonfactively represented by a sentence. Perhaps there is, for instance, no fact of the matter which worldly proposition is represented by the sentence "The mass of the ball is one kilogram" in a Newtonian language. If so, there is some truth in approach (a) to the general issue that a theory may not uniquely determine representation relations, namely the approach that says that such a unique determination is not necessary in order for a theory of representation to be satisfactory. This truth applies more widely than just in the cases mentioned above. So, the fact that our account of representation does not always yield a unique mapping between representings and representeds may ultimately be an advantage rather than a problem.

To sum up, our account of representation can easily handle misrepresentation, it can account for the asymmetry of the representation relation (despite taking representation to be a matter of sameness of form), and it has the resources to address the potential problem that there may be several distinct isomorphisms between normative and alethic incompatibilities.

4.1.4 Reason Relations for Representationalists

According to a semantics-first approach to reason relations, we should explain reason relations among sentences in terms of what these sentences represent. What could such an account look like?

In the previous chapters, we offered an account of reason relations in a normative-pragmatic metavocabulary. According to this account, a sentence is a reason for another sentence if asserting the first and denying the second is normatively impossible (or out-of-bounds, in bilateralist terminology) because one cannot be entitled to such a combination of discursive acts. A sentence is a reason against another sentence if asserting both sentences is normatively forbidden. Given the isomorphism that underwrites the sharing of rational forms, we can give a parallel account of reason relations in terms of alethic modal relations among states, namely the following: A sentence is a reason for another sentence if and only if the combination of any truth-maker of the worldly proposition represented by the first sentence and any falsity-maker of the worldly proposition represented by the second sentence is impossible, in the alethic modal sense. And a sentence is a reason against another sentence if and only if the combination of truth-makers of the two worldly propositions represented by the sentences is alethically impossible.

On this view, a sentence is a reason for another sentence if the truth of the first rules out the falsity of the second. And it is a reason against the other sentence if they cannot be true together. Such impossibilities hold because of the modal relations among the states that make up the worldly propositions that the sentences represent. Thus, reason relations among sentences are explained in terms of relations among what they represent.

At this point, a reminder of what we mean by a sentence being a reason for or against another sentence is in order. As in the previous chapters, the reason relations at issue here do not determine but merely constrain the "reasons" (in the more common sense) we have for particular beliefs in particular circumstances. We are defining reason relations of implication and incompatibility, which are at one remove from practical policies. That one sentence is a reason for another does not imply that you ought to believe what the second sentence says if you believe what the first sentence says. Rather, when you believe what a sentence says on the grounds of what another sentence says, then you believe what the first sentence says for good reason only if the second sentence is a reason for the first sentence, in the sense at issue in this book. And the same holds *mutatis mutandis* for reasons against something.

4.2 Truth-Maker Theory

In the first section of this chapter, we have outlined in very abstract terms the strategy of a particular semantics-first approach to content and reason relations. We have already seen that the idea of representation seems to lead us back to the pragmatics-first approach from the previous chapters. In particular we seem compelled to talk about the norms that govern the use of sentences when we want to say why a particular sentence represents what it represents. In order to make such connections precise, however, we need a formal codification of the representational account sketched above. Fortunately, there is an extant formal theory that can serve this purpose with only minor adjustments, namely the exact truth-maker theory developed by Kit Fine (2017a; 2017b). The goal of this section is to put in place a version of truth-maker theory, which we will later show to be isomorphic to the theory that we have put forward in the previous chapters.

Due to this isomorphism, the ability of the sequent calculus NMMS from the previous chapter to codify open reason relations will extend to this version of truth-maker theory. Thus, the truth-maker semantic theory as amended below is substructural in the same way in which the sequent calculus NMMS is substructural. In order to meet our desideratum of being able to codify open reason relations, we will start with a structural version of truth-maker theory and then show how we can drop the constraints that prevent the closure-structured version from codifying open reason relations.

Truth-maker theory is perhaps the most powerful and flexible version of the broad idea of truth-conditional semantics. Truth-maker theory has been used in metaphysics, for instance, to give accounts of essence (Zylstra, 2019; Hale, 2020), grounding (deRosset and Fine, 2022), and generalized identity (Elgin, 2021). It has been applied in natural language semantics, for instance, to give accounts of attitude reports and deontic modals (Moltmann, 2020), conditionals (Yablo, 2016; Santorio, 2018), and to give accounts of the subject matter of sentences (Yablo, 2014; Fine, 2020). Moreover, truth-maker theory has been used to give accounts of various normative phenomena (Fine, 2018c, a, b), verisimilitude (Fine, 2019), truth (Jago, 2018), and other phenomena. One reason why accounts of all these different and philosophically puzzling phenomena can be formulated in truth-maker theory is that it is hyperintensional. That is, we can distinguish between the contents of sentences whose truth-values are necessarily the same, including sentences that are equivalent according to classical logic. Roughly, the idea is that, say, "$1 + 1 = 2$" and "π is irrational" might differ in content, although they are both necessarily true, because they are made true by different aspects of reality, or states. Similarly, "It is raining" and "It is raining and there is or isn't extraterrestrial life" differ in content,

although they are equivalent according to classical logic, because there are aspects of reality, or states, that are relevant to the truth of the second but that are not relevant to the truth of the first, namely how things stand with extraterrestrial life. Moreover, truth-maker theory has the resources to assign distinct semantic values to two necessarily false sentences, such as "It is day and also night" and "The ball is blue all over and yellow all over," namely by allowing for distinct impossible states, such as a state of it being day and night and a distinct state of the ball being blue all over and yellow all over. Thus, we can make very fine-grained distinctions in truth-maker theory, and it is in this sense much more expressively powerful than, for instance, possible world semantics. This is part of what makes this framework flexible enough to allow us to formulate an account of reason relations that is isomorphic to the one from the previous chapters.

Let us formulate in more detail the precise version of truth-maker theory that is relevant for our project. The starting point of truth-maker theory is the idea that there are states that can obtain or not obtain and that make some sentences true, that is, verify them, and make some sentences false, that is, falsify them (Fine, 2017a, 2014). The verifying states (truth-makers) and falsifying states (falsity-makers) of a sentence are wholly relevant to the sentence. The state of it raining, for instance, is a truth-maker of the sentence "It is raining." But the state of it raining and it being cold is not a truth-maker of "It is raining" because it is not wholly relevant to the truth of the sentence (it's being cold is irrelevant). States can be parts of other states. Some states are possible and others are impossible.

Let us follow Fine (2017a, 647) and say that a "modalized state space" is a triple of a set of such states, the subset of those states that are possible states, and the parthood relation among these states:

Definition 35 (Modalized state space)**.** A modalized state space is a triple, $\langle S, S^{\diamond}, \sqsubseteq \rangle$, such that S is a non-empty set of states, $S^{\diamond} \subseteq S$ (intuitively: the possible states), and $\sqsubseteq$ is a partial order on S (intuitively: parthood), such that all subsets of S have a least upper bound.

States can be combined—or, as we shall say, "fused"—to yield further states. We define the fusion of states as the smallest state of which all the initial states are parts, that is, as the least upper bound of the initial set of states under our partial order (Fine, 2017a, 646).

Definition 36 (Fusion)**.** The fusion of a set of states $T = \{t_1, t_2, t_3, ...\}$, written $t_1 \Cup t_2 \Cup t_3 ...$ or $\Cup T$, is the least upper bound of T with respect to $\sqsubseteq$.[17]

Every modalized state space has a least element, $\square$, which is the "null state" that is part of every state. It is the least upper bound of the empty set.

We said above that a sentence, ϕ, of language $\mathfrak{L}$ represents a worldly proposition, which is a pair, $\langle |\phi|^+, |\phi|^- \rangle$, of the set of the sentence's verifiers, $|\phi|^+$, and the set of its falsifiers, $|\phi|^-$. We call a mapping of sentences to worldly propositions an interpretation function. A model is a modalized state space together with an interpretation function.

Definition 37 (Model)**.** Given a language $\mathfrak{L}$, a model, $\mathcal{M}$, is a quadruple $\langle S, S^\diamond, \sqsubseteq, |\cdot| \rangle$, where $\langle S, S^\diamond, \sqsubseteq \rangle$ is a modalized state space and $|\cdot|$ is a interpretation function, such that $|A| = \left\langle |A|^+, |A|^- \right\rangle \in \mathcal{P}(S) \times \mathcal{P}(S)$.

We write $\mathcal{M}, s \Vdash A$ if a state s verifies a sentence A in model $\mathcal{M}$, and if no risk of confusion arises, simply $s \Vdash A$. Similarly, we write $\mathcal{M}, s \dashv\!\Vdash A$ to say that s is a falsifier of sentence A in model $\mathcal{M}$. In other words, relative to a given model, $|A|^+$ is the same as $\{s : s \Vdash A\}$, and $|A|^-$ is the same as $\{s : s \dashv\!\Vdash A\}$.

In order to give a treatment of logical vocabulary in terms of these models, we have to treat some bits of vocabulary as logical vocabulary by keeping its meaning fixed in all models. So let $\mathfrak{L}$ be a language that results from adding $\neg$ and $\wedge$ and $\vee$ to a stock of atomic sentences $\mathfrak{L}_{\mathfrak{B}}$. To hold the meanings of logical vocabulary fixed, we give semantic clauses that specify the truth-makers and falsity-makers of complex sentences in terms of the truth-makers and falsity-makers of their parts.[18] We will limit ourselves throughout to propositional logic, as we did in the previous chapters.[19] For the atomic sentences, our models directly specify their truth-makers and falsity-makers. Thus, we have:

(atom+) $s \Vdash p$ if and only if $s \in |p|^+$

(atom−) $s \dashv\!\Vdash p$ if and only if $s \in |p|^-$

For conjunction, it is plausible that the truth-makers of conjunctions are fusions of the truth-makers of their conjuncts. For example, the sentence "It is raining and it is night" is made true by states that combine a truth-maker of "It is raining" and a truth-maker of "It is night" (and nothing else). And it also seems plausible that a conjunction is made false by any state that makes one of the conjuncts false, or a fusion of such states. For example, the state of it being sunny and the state of it being daytime are intuitively both falsity-makers for "It is raining and it is night." Thus, we have the following semantic clauses for conjunction[20]:

(and+) $s \Vdash B \wedge C$ if and only if
$\exists u, t\,(u \Vdash B$ and $t \Vdash C$ and $s = u \Cup t)$

(and−) $s \dashv\!\Vert B \wedge C$ if and only if
$s \dashv\!\Vert B$ or $s \dashv\!\Vert C$ or $\exists u, t\,(u \dashv\!\Vert B$ and $t \dashv\!\Vert C$ and $s = u \Cup t)$

Disjunction works, plausibly, in an analogous way, with truth-making flipped to falsity-making and vice versa: A state makes a disjunction true if it makes one of the disjuncts true or is a fusion of such states. And this suggests the following clause for the truth-makers of disjunctions.

(or+) $s \Vdash B \vee C$ if and only if
$s \Vdash B$ or $s \Vdash C$ or $\exists u, t\,(u \Vdash B$ and $t \Vdash C$ and $s = u \Cup t)$

Moreover, a state makes a disjunction false if it makes both disjuncts false.

(or−) $s \dashv\!\Vert B \vee C$ if and only if
$\exists u, t\,(u \dashv\!\Vert B$ and $t \dashv\!\Vert C$ and $s = u \Cup t)$

And for negation, it is plausible to assume that a state makes a negation true if it makes the negatum false, and it makes a negation false if it makes the negatum true.

(neg+) $s \Vdash \neg B$ if and only if $s \dashv\!\Vert B$

(neg−) $s \dashv\!\Vert \neg B$ if and only if $s \Vdash B$

It will prove convenient below not only to talk of individual sentences being made true or false, but also of sets of sentences being made true or false. We will say that a set of sentences is made true by the truth-makers of the conjunctions of its members, unless there are no such truth-makers, in which case it is made true by the null state. Similarly, we say that a set of sentences is made false by the falsity-makers of the disjunction of its members, unless there are no such falsity-makers.[21]

Definition 38 (Truth- and falsity-makers of sets). $u \Vdash \Gamma$ if and only if $u \Vdash \bigwedge \Gamma$, unless $\{x : x \Vdash \bigwedge \Gamma\} = \emptyset$ in which case $\square$ and nothing else makes Γ true. And $t \dashv\!\Vert \Delta$ if and only if $t \dashv\!\Vert \bigvee \Delta$, unless $\{x : x \dashv\!\Vert \bigvee \Delta\} = \emptyset$ in which case $\square$ and nothing else makes Δ false.

Given the clauses for conjunction and disjunction, it is easy to see that the truth-makers of sets are the fusions of the truth-makers of their members, if there are any. And the falsity-makers of sets are the fusions of the falsity-makers of their members, if there are any. So the appeal to conjunction and disjunction could be eliminated; it is merely an efficient way to talk about

fusions. The unless-clauses are useful because they imply that the empty set is made true and false by the null state. This might seem strange but it has technical advantages. In particular, it allows us to say that fusing a state with the truth-maker or falsity-maker of the empty set simply returns the original state; for, $s \Cup \square = s$.

It will also occasionally be useful to allow ourselves to ignore certain cardinality issues. To see what we mean, note that, given a countable infinity of states, there will be an uncountable infinity of worldly propositions. Hence, we cannot express all worldly propositions in a language. It is sometimes useful to ignore this complication, and we can do so by making the following assumption.[22]

Assumption 39 (*Expressibility of worldly propositions*)**.** For any worldly proposition $\langle V, F\rangle$, we can add sentences, $\Gamma \cup \Delta$, to our language such that V contains exactly one truth-maker for each sentence in Γ (and nothing else) and F contains exactly one falsity-maker for each sentence in Δ (and nothing else).

With this assumption in place, if we encounter a worldly proposition, we can assume that our language has (or can be given) the resources to express the proposition, and to do so in a sufficiently fine-grained way, namely as fine-grained as the states in the proposition. As will become clear in the next two chapters, the assumption is best understood as the mirror image of the fact that our grasp of states is limited by our language.

We have now spelled out most of the details of our models. However, one crucial element is missing. So far, we have put no constraints on the distinguished set of possible states; they might be any subset of our states. Fine often imposes the following constraints, and they will become important below:

Downward-Closure: If $s \in S^{\diamond}$ and $t \sqsubseteq s$, then $t \in S^{\diamond}$.

Exclusivity: If $s \in |p|^{+}$ and $t \in |p|^{-}$, then $\forall u(s \Cup t \Cup u \notin S^{\diamond})$.[23]

Exhaustivity: $\forall u \in S^{\diamond}$, either $\exists s \in |p|^{+}\ (u \Cup s \in S^{\diamond})$ or $\exists t \in |p|^{-}\ (u \Cup t \in S^{\diamond})$.

Downward-Closure says that all parts of a possible state are possible. Exclusivity says that if you take any atomic[24] sentence, then if you fuse one of its verifiers with one of its falsifiers together with any other state (or, in effect, none, if it is the null state), you always get an impossible state. And Exhaustivity says that if you have a possible state and an atomic sentence, then you can extend it to a possible state either by fusing it with a verifier of the sentence or by fusing it with a falsifier of the sentence.

It will emerge in due course that the three constraints of Downward-Closure, Exclusivity, and Exhaustivity correspond, respectively, to the structural rules of Weakening, Containment, and Cut in a classical sequent calculus. Moreover, the semantic clauses correspond to our operational sequent rules. This correspondence is at the center of the isomorphism we will establish between the theory from the previous chapters and truth-maker theory.

Before we get to this correspondence, however, we must address what is perhaps the central question in formulating any logical system: How should we define consequence? (For us, this shows up as the key *semantic* question about a logic.) It is common in the literature to define consequence in truth-maker theory in non-modal terms, that is, without any appeal to the distinction between possible and impossible states. In particular, Fine often uses the following notions of consequence[25]:

- *Entailment*: P entails Q if and only if every verifier of P is a verifier of Q (Fine, 2017a, 640–41).
- *Containment*: Q contains P if and only if (i) every verifier of Q includes as a part a verifier of P and (ii) every verifier of P is included as a part in a verifier of Q (Fine, 2017a, 640–41).

These notions can be extended to the case of multiple premises in different ways (Fine and Jago, 2019). As we will see below, however, an interesting equivalence emerges between truth-maker theory and the theory in the previous chapters if we adopt, instead, a different notion of consequence, which we call "truth-maker consequence" or "TM-validity"[26] and write as $\models_{TM}$. This notion captures, within truth-maker theory, the representationalist account of reason relations sketched in the previous section.

Recall that we suggested above that, from a representationalist perspective, what it means for A to be a reason for B is that it is impossible that A is true and B is false. And what it means for A to be a reason against B is that it is impossible that A is true and B is also true. We can now formulate this idea in truth-maker theory as follows: A is a reason for B if and only if any fusion of a truth-maker of A and a falsity-maker of B is an impossible state. And A is a reason against B if and only if any fusion of a truth-maker of A and a truth-maker of B is an impossible state. As in previous chapters, we can think of these relations as codified in a consequence relation by saying that B is a consequence of A if and only if A is a reason for B, and the empty set is a consequence of A together with B if and only if A is a reason against B. Notice that this truth-maker conception of consequence, unlike Fine's candidates, is a *modal* notion; it crucially involves the distinction between possible and impossible states.

According to this notion, B is a consequence of A just in case no truth-maker of A is compossible with any falsity-maker of B.[27] If we finally also generalize these intuitive ideas to sets of sentences, the result is the following:

> *Truth-Maker Consequence*: $\Gamma \models_{TM} \Delta$ (in a model) if and only if (in that model) any fusion of verifiers for each $\gamma \in \Gamma$ and falsifiers for each $\delta \in \Delta$ is an impossible state, that is, $s \notin S^{\diamond}$ for all $s = u \Cup t$ such that $u \Vdash \Gamma$ and $t \dashv\!\Vdash \Delta$.[28]

This notion of consequence is essentially both modal and bilateralist, in that it makes crucial use of the distinction between possible and impossible states and, moreover, makes crucial use of falsity-makers as well as truth-makers. To the best of our knowledge, these features distinguish this notion of consequence from all extant notions of consequence used in truth-maker theory.[29] As we shall see below, this allows us to make full use of the many fine-grained features of truth-maker theory in thinking about consequence. In particular, it will allow us to capture open reason relations.

It will prove useful in the next chapter to add a definition of an implication relation among worldly propositions, that is, a kind of consequence relation that does not depend on any particular language. We can apply the idea behind Truth-Maker Consequence directly to worldly propositions by saying that a set of worldly propositions implies another set of worldly propositions just in case any fusion of truth-makers for each proposition in the first set and falsity-makers for each proposition in the second set is an impossible state.

> *Propositional Implication*: With G and D being sets of worldly propositions, $G \Vdash_{PI} D$ (in a modalized state space) if and only if (in that modalized state space) any fusion of verifiers for each $\langle V, F\rangle \in G$ and falsifiers for each $\langle V, F\rangle \in D$ is an impossible state.

We could think of the worldly propositions that would be the interpretation of a negation of a sentence as the (propositional) negation of the proposition that is the interpretation of the negated sentence, and similarly for the other connectives. For example, we can think of $\langle |\phi|^{-}, |\phi|^{+}\rangle$ as the (propositional) negation of $\langle |\phi|^{+}, |\phi|^{-}\rangle$. While this way of thinking will become relevant in the next chapter, however, it is more illuminating for the purposes of the current chapter to think of consequence as a relation among sets of sentences of a language.

Truth-Maker Consequence gives us a notion of consequence relative to models. And if we have a model that models reality (including modal aspects of reality) in the right way, it will be consequence relative to that

model that captures what follows from what in the right way. In specifying a particular model, we can choose which material inferences are TM-valid in that model. If we want to ensure, for instance, that "*o* is colored" follows from "*o* is green," all we have to do is to ensure that every fusion of a falsity-maker for "*o* is colored" and a truth-maker for "*o* is green" is an impossible state.[30] Thus, truth-maker consequence is not logical consequence; it is not closed under uniform substitution of non-logical expressions. We are, of course, especially interested in these wider notions of consequence. We can, however, easily restrict truth-maker consequence to logical consequence (LTM-validity) by quantifying over models in the usual way.

> *Logical Truth-Maker Consequence*: $\Gamma \underset{LTM}{\models\!\!=} \Delta$ if and only if, in all models, $s \notin S^{\diamond}$ for all $s = u \Cup t$ such that $u \Vdash \Gamma$ and $t \dashv\!\!\Vdash \Delta$.

As we will see shortly, making different choices in the semantic set-up leads to different consequence relations for TM-validity and LTM-validity. Hence, these are really families of consequence relations. We will disambiguate where necessary by using appropriate labels. It will turn out that different members of these families correspond exactly to the different accounts of consequence from the previous chapter.

4.3 Articulating the Isomorphism

Let us now bring out how the version of truth-maker theory maps onto our sequent calculus NMMS from the previous chapter. As we will see, the left-rules of NMMS turn out to correspond to the semantic clauses for truth-makers. The right-rules turn out to correspond to the clauses for falsity-makers. And the structural rules of Containment, Monotonicity, and Cut will turn out to correspond to constraints on possible states, namely Exclusivity, Downward Closure, and Exhaustivity, respectively. In this section, we will explain these correspondences and some of their implications.[31] In order to do so, we start by explaining the general idea or strategy behind the isomorphism. We then turn to the operational sequent rules, and we end by discussing the structural rules.

4.3.1 The Idea Behind the Isomorphism

Recall that according to the pragmatics-first approach that we have presented in the previous chapters, what it means to say that Δ follows from Γ is that if one is committed to assert everything in Γ, then one cannot be entitled to deny everything in Δ. Equivalently, we said that what it means that Δ follows from Γ is that any combination of assertions of every sentence in Γ and denial of every sentence in Δ is

normatively impossible (forbidden) or "out-of-bounds." Let us denote this notion of consequence by $\models_{PN}$, for Pragmatic-Normative consequence. This pragmatic-normative understanding of consequence offers an intuitive interpretation of sequents and sequent rules. In particular, sequent rules can then be interpreted as telling us that certain combinations of assertions and denials are normatively impossible if other such combinations are normatively impossible.

In formulating Sharing-Rational-Forms above, we explained that a sentence and a worldly proposition share their rational forms if their occurrences stand in isomorphic modal relations to other occurrences of sentences and worldly propositions respectively. On the side of sentences, the relevant modality is normative and the occurrences are assertions and denials. And on the side of worldly propositions, the relevant modality is alethic and the occurrences are truth-makers and falsity-makers.

We can now see what it would mean to establish the required isomorphism between sentences and worldly propositions in terms of pragmatic-normative consequence and truth-maker consequence. There would have to be an interpretation that maps sentences to worldly propositions in such a way that pragmatic-normative consequence and truth-maker consequence coincide. We can now formulate this idea as follows:

Sharing-Rational-Forms (formal)
A sentence, A, of a language and the worldly proposition, $|A|$, that is its interpretation (in a model) share their rational form if and only if for every two collections of sentences Γ and Δ,
(pos) $\Gamma, A \models_{PN} \Delta$ if and only if $\Gamma, A \models_{TM} \Delta$
(neg) $\Gamma \models_{PN} A, \Delta$ if and only if $\Gamma \models_{TM} A, \Delta$.

This says that a sentence A shares its rational form with its interpretation $|A|$ (in a model) just in case an assertion of A is normatively incompatible with a collection of assertions and denials if and only if (in that model) every fusion of a truth-maker of A with truth-makers for the assertions and falsity-makers for the denials is an impossible state. And *mutatis mutandis* for denials of A. Thus, this is really just another formulation of Sharing-Rational-Forms from Section 4.1 above.

We have interpreted sequent calculi in a pragmatic-normative way in the previous chapters. And we have introduced truth-maker consequence in the previous section. So, we can now establish the desired isomorphism by mapping truth-maker consequence into our sequent calculus. Crucially, however, this mapping is not just global but also local, that is, we will map the details of the two frameworks to each other, and this will allow us to

tweak both frameworks in exactly parallel ways, including the rejection of global closure-structural principles.

To see how this mapping of details works, notice the common structure of pragmatic-normative consequence and truth-maker consequence: First, we have items with positive and negative occurrences. Second, we have a notion of combining these elements into larger structures. Third, we distinguish among such combinations between the possible and the impossible ones. We can read sequent rules as concerning these three elements, and we can, thus, generate two parallel readings of every sequent rule, one pragmatic-normative reading and one truth-maker reading. Take, for instance, the following version of the structural rule of Cut.

$$\frac{\Gamma \succ \Delta, A \qquad A, \Gamma \succ \Delta}{\Gamma \succ \Delta}\ [\text{Cut}]$$

Reading this rule contrapositively, it says that, for any Γ and Δ, if the sequent $\Gamma \succ \Delta$ doesn't hold, then either $\Gamma \succ \Delta, A$ or $A, \Gamma \succ \Delta$ doesn't hold. We can formulate the pragmatic-normative (PN) and the truth-maker bilateralist (TM) readings of this claim as follows:

PN-Cut: For any normatively possible combinations of assertions and denials and any sentence, A, one can extend the combination while keeping it normatively possible either by asserting or by denying A.

TM-Cut: For any alethically possible state and any worldly proposition, A, one can extend the state into an alethically possible state by fusing it with either a verifier or a falsifier of A.

Notice that TM-Cut is a version of Fine's Exhaustivity, from the previous section. The only two apparent differences are that Exhaustivity is formulated without mentioning worldly propositions and that it concerns only atomic sentences. Both differences are, however, insubstantial. Below we will spell this equivalence of Exhaustivity and Cut out in a more rigorous fashion.

An analogous result holds for all of the other rules. To see this, let's consider next the principle of Containment, which says that every sequent of the form $\Gamma, A \succ A, \Delta$ holds. The two readings are the following:

PN-CO: Any combination of assertions and denials in which any sentence is asserted and also denied is normatively impossible.

TM-CO: Any state that includes a verifier and also a falsifier for any worldly proposition is alethically impossible.

Notice that TM-CO is a variant of Exclusivity from the previous section. The only and irrelevant difference is that Exclusivity is formulated only for atomic sentences.

The third constraint on possible states, which is Downward Closure, is equivalent to the structural rule of [Weakening] or Monotonicity, which says that if $\Gamma \succ \Delta$ holds, then so does $\Gamma, \Theta \succ \Delta, \Sigma$. The corresponding two readings are these:

PN-MO: All combinations of assertions and denials that include a normatively impossible combination as a part are themselves normatively impossible.

TM-MO: All states that include an alethically impossible state are themselves alethically impossible.

We can also give two interpretations of each of our operational rules. The left-rules and the right-rules all work generically in the same way. Hence, we can formulate the equivalence at this generic level. For the left-rules, we can provide parallel interpretations in the following way:

PN-left: Left-rules specify the contributions that the assertions of complex sentences make to collections of assertions and denials being normatively impossible in terms of the contributions made by the assertions or denials of their constituent sentences.

TM-left: Left-rules specify the contributions that the verifiers of the worldly proposition represented by complex sentences make to states being alethically impossible in terms of the contributions made by the verifiers or falsifiers of their constituent sentences.

And the parallel formulations for the right-rules are as follows:

PN-right: Right-rules specify the contributions that the denials of complex sentences make to collections of assertions and denials being normatively impossible in terms of the contributions made by the assertions or denials of their constituent sentences.

TM-right: Right-rules specify the contributions that the falsifiers of the worldly proposition represented by complex sentences make to states being impossible in terms of the contributions made by the verifiers or falsifiers of their constituent sentences.

We thus have two parallel readings of valid sequents and sequent rules: one alethic-modal, the other pragmatic-normative. In both interpretations, the consequence relation holds between two sets just in case certain states or combinations of assertions and denials are impossible. This impossibility is normative in one case and alethic in the other case.

According to the account of representation from the first section of this chapter, a sentence and a worldly proposition share their rational form just in case the incompatibilities of assertions and denials of the sentence with other assertions and denials are isomorphic to the incompatibilities of the truth-makers and falsity-makers of the worldly proposition with other states. In order for a sentence to represent a worldly proposition, the two must share their rational form.[32] So if we want our sentences to represent particular worldly propositions, we must make it the case that their assertions and denials are governed by norms that are isomorphic in this way. The two readings of sequent rules just sketched means that if such an isomorphism holds for our atomic sentences, then it will also hold for complex sentences, given that we use corresponding sequent rules and truth-maker clauses for our logical vocabulary.

4.3.2 Operational Rules

In this and the following subsection we want to spell out the isomorphism just sketched with more formal rigor. We start with the operational rules in this subsection and turn to the structural rules in the next subsection.

Before we discuss particular sequent rules in detail, however, it will prove useful to introduce the notion of a sequent "deeming certain states impossible." The idea is that a sequent deems exactly those states impossible that need to be impossible in order for the sequent to hold, that is, for the conclusions to be a consequence of the premises. Hence, a sequent deems impossible exactly those states that are fusions of truth-makers of the premises and falsity-makers of the conclusions.

Definition 40 (Deeming impossible)**.** A sequent $\Gamma \succ \Delta$ deems all and only those states impossible that are a fusion of verifiers for everything in Γ and falsifiers for everything in Δ, that is, any state $s = u \Cup t$ such that $u \Vdash \Gamma$ and $t \dashv\!\| \Delta$.

This notion will prove useful because it turns out that, in all our sequent rules, the set of states that are deemed impossible by the bottom-sequent is the union of the sets of states deemed impossible by the top-sequents.

Let us start our discussion of the operational rules by looking at the NMMS left-rule for conjunction, which we repeat here for convenience:

$$\frac{\Gamma, A, B \succ \Delta}{\Gamma, A \wedge B \succ \Delta}\ [\mathrm{L}\wedge]$$

Recall that, like all of our operational rules, this rule is invertible, so that it holds not only from top to bottom but also from bottom to top. Hence, [L∧] implies that $\Gamma, A, B \succ \Delta$ is derivable if and only if $\Gamma, A \wedge B \succ \Delta$ is derivable. By the definition of truth-maker consequence, this rule is sound for truth-maker consequence just in case, in all models, if we take any state, $g \Cup d$, that is a fusion of truth-makers for each member of Γ and falsity-makers for each member of Δ, then all fusions of $g \Cup d$ with a truth-maker, a, of A and a truth-maker, b, of B are impossible if and only if all fusions of $g \Cup d$ with any truth-maker, ab, of $A \wedge B$ are impossible. Since this must hold for arbitrary states $g \Cup d$, it includes the case where $g \Cup d = \square$. Hence, it must be that $\square \Cup a \Cup b$ is impossible if and only if $\square \Cup ab$ is impossible and, so, $a \Cup b$ is impossible if and only if ab is impossible.

Now, recall the semantic clause for the truth-makers of conjunctions from above, which we repeat for convenience with the variables adjusted to match our present case.

(and+) $\quad ab \Vdash A \wedge B$ if and only if
$\quad\quad \exists a, b\,(a \Vdash A$ and $b \Vdash B$ and $ab = a \Cup b)$

This clause says that the truth-makers of $A \wedge B$ are all and only the fusions of a truth-maker of A and a truth-maker of B. Thus, the states that our sequent rule requires to be co-impossible in all models are precisely those states that our semantic clause requires to be identical in all models. Because these states are identical, the results of fusing them with $g \Cup d$ are also identical. So, our semantic clause (and+) is sufficient to ensure that [L∧] is valid for truth-maker consequence.

Moreover, our semantic clause (and+) is also necessary for the validity of [L∧]. For, suppose that there were a state ab that is a truth-maker of $A \wedge B$ but is not a fusion of a truth-maker of A and a truth-maker of B. We could then find a model in which any fusion of the form $g \Cup d \Cup ab$ is impossible while fusing $g \Cup d$ with some truth-maker of A and some truth-maker of B is possible. So, in that model, $\Gamma, A, B \not\models_{TM} \Delta$ but $\Gamma, A \wedge B \models_{TM} \Delta$. And we can generalize this strategy to cover all models, so that we have a counterexample to [L∧].

What this means is that our semantic clause for the truth-makers of conjunctions, namely (and+), and our left-rule for conjunction in the sequent calculus, namely [L∧], have the same impact on the consequence relation. They are equivalent in their significance for the consequence

relations defined in the two frameworks. Another way to look at this fact is to note that the sequents $\Gamma, A, B \succ \Delta$ and $\Gamma, A \wedge B \succ \Delta$ deem exactly the same states impossible. A corresponding result holds for all of our operational rules, namely:

Lemma 41. *For every application of an operational rule of* NMMS, *the set of states deemed impossible by the bottom-sequent is the union of the sets of states deemed impossible by the top-sequents. (Appendix, Lemma 50)*

The reason why this holds is that the states that the active sentence in the bottom-sequent contributes to the fusions that the bottom-sequent deems impossible are exactly the states that the active sentences in the top-sequents contribute to their fusions.

To see what we mean, consider the NMMS right-rule for conjunction as another example. Recall from the previous chapter that the NMMS right-rule for conjunction is the following:

$$\frac{\Gamma \succ \Delta, A \qquad \Gamma \succ \Delta, B \qquad \Gamma \succ \Delta, A, B}{\Gamma \succ \Delta, A \wedge B}\ [\mathrm{R}\wedge]$$

And recall that the semantic clause for falsity-makers of conjunctions from above is the following, repeated here for convenience:

(and−) $s \dashv\!\| B \wedge C$ if and only if
$s \dashv\!\| B$ or $s \dashv\!\| C$ or $\exists u, t\,(u \dashv\!\| B$ and $t \dashv\!\| C$ and $s = u \Cup t)$

Note that the bottom-sequent, $\Gamma \succ A \wedge B, \Delta$, of our right-rule holds, by the definition of truth-maker consequence, just in case every fusion of truth-makers for Γ and falsity-makers for Δ with any falsity-maker of $A \wedge B$ is an impossible state. Now, by our semantic clauses, a falsity-maker of $A \wedge B$ is a state that is either a falsity-maker of A or a falsity-maker of B, or a fusion of such falsity-makers. But the fusions of those states with truth-makers for Γ and falsity-makers for Δ are precisely the states that are deemed impossible by the top-sequents of [R∧]. So, the set of states deemed impossible by the bottom sequent, $\Gamma \succ A \wedge B, \Delta$, is precisely the union of the states deemed impossible by the top-sequents of [R∧] according to (and−).

Similarly, for negation, since a truth-maker of A is a falsity-maker of $\neg A$, the sequents $\Gamma, A \succ \Delta$ and $\Gamma \succ \neg A, \Delta$ deem exactly the same states impossible. And since a falsity-maker of A is a truth-maker of $\neg A$, the sequents $\Gamma \succ A, \Delta$ and $\Gamma, \neg A \succ \Delta$ deem exactly the same states impossible. That corresponds to the right-rule and the left-rule for negation.

The lemma above says that this is always the case, that is, it holds not only for conjunction and negation but also for disjunction. It follows that

applications of our sequent-rules are guaranteed to preserve truth-maker consequence.

Proposition 42. *All the operational rules of* NMMS *are sound for* $\models_{TM}$, *that is, if all the top-sequents of a rule-application are TM-valid, then so is the bottom sequent. (Appendix, Proposition 51)*

Indeed, our lemma shows not only that our sequent rules are sound but also that they are complete with respect to the corresponding truth-maker clauses. To see the outlines of an argument for why this holds (see the Appendix for more details), suppose that a sequent is not derivable. Then for any rule-application that would conclude that sequent, one of the top-sequents must be underivable. And when we trace this underivability up a potential proof-tree, we will reach a sequent whose underivability depends merely on atomic sentences. It can be shown that we can construct a counterexample for such a sequent by specifying a model in which a fusion of truth-makers for its premises and falsity-makers for its conclusions is a possible state. And thanks to our lemma, the underivable sequent with which we started must deem that state in that model impossible. Hence, we have a counterexample to our underivable sequent.[33] So the sequent rules are both sound and complete with respect to the truth-maker semantics.

4.3.3 Structural Rules

This brings us to the structural rules. We won't consider permutation and contraction, but continue to think in terms of sets of sentences. Thus, we have just three structural rules: [Containment], [Weakening], and [Cut]. These are the sequent rules that correspond to (versions of) the reflexivity, monotonicity, and transitivity of consequence. They are the constraints that allow our consequence relations to encode only topologically closed reason relations. If we want to be able to theorize open reason relations, we must relax these constraints. Hence, if we want to formulate versions of truth-maker theory that can capture open reason relations, it is crucial for us to understand what corresponds to these structural constraints in truth-maker theory and how we can relax them. We repeat the structural sequent rules here for convenience:

$$\frac{}{\Gamma, A \succ A, \Delta}\ [\text{Containment}]$$

$$\frac{\Gamma \succ \Delta}{\Gamma, \Theta \succ \Delta, \Lambda}\ [\text{Weakening}] \qquad \frac{\Gamma \succ \Delta, A \qquad A, \Gamma \succ \Delta}{\Gamma \succ \Delta}\ [\text{Cut}]$$

As already intimated, the three structural rules correspond exactly to the three constraints on possible states. Exclusivity corresponds to [Containment]. Downward-Closure corresponds to [Weakening]. And Exhaustivity corresponds to [Cut]. Let us run through the arguments for the three rules.

Recall that Exclusivity says that if $s \in |p|^+$ and $t \in |p|^-$, then $\forall u(s \Cup t \Cup u \notin S^\Diamond)$, that is, if you fuse a truth-maker and a falsity-maker of an atomic sentence with any other state, then the result is always an impossible state. If that holds in a model, then clearly $\Gamma, p \models_{TM} p, \Delta$ in that model. And if it holds in all models, then $\Gamma, p \models_{LTM} p, \Delta$. But that is precisely what the [Containment] rule allows us to prove; for it allows us to derive $\Gamma, p \succ p, \Delta$. Moreover, for the other direction, in order to ensure that all the sequents of the form $\Gamma, p \succ p, \Delta$ hold according to the logic of truth-maker consequence, we must accept Exclusivity as a constraint on possible states.

Downward-Closure says that if $s \in S^\Diamond$ and $t \sqsubseteq s$, then $t \in S^\Diamond$, that is, if a state is possible, then all of its parts are possible. So, if a state is impossible, then all the states of which it is a part are impossible. A sequent $\Gamma \succ \Delta$ is TM-valid, in a model, just in case all the states that are fusions of truth-makers of Γ and falsity-makers of Δ are impossible. If that is the case and Downward-Closure holds, then all the states that include such a fusion as a part are also impossible. Now, such a fusion is part of every state that is deemed impossible by a sequent of the form $\Theta, \Gamma \succ \Delta, \Sigma$. For these states are fusions of truth-makers of Γ and falsity-makers of Δ with some further states, namely truth-makers of Θ and falsity-makers of Σ. Hence, if $\Gamma \succ \Delta$ is TM-valid, then so is $\Theta, \Gamma \succ \Delta, \Sigma$. Thus, Downward-Closure is sufficient for the soundness of [Weakening].

Downward-Closure is also necessary for the soundness of [Weakening] if we assume the expressibility of propositions. For suppose that Downward-Closure fails, so that there is an impossible state, s, that is part of a possible state, $s \Cup t$. If we can find a sequent, $S \succ$, that deems only the impossible state s, impossible, and also a sentence, T, whose only truth-maker is t, then $S \models_{TM}$ but $S, T \not\models_{TM}$. And that gives us a counterexample to [Weakening]. Hence, Downward-Closure and [Weakening] are equivalent regarding their effects on the consequence relation in their respective frameworks.

Lastly, remember that Exhaustivity says that if a state, u, is possible, then either $\exists s \in |p|^+ \ (u \Cup s \in S^\Diamond)$ or $\exists t \in |p|^- \ (u \Cup t \in S^\Diamond)$. That is, for any atomic sentence, a possible state can always be extended to another possible state by fusing it either with a truth-maker or with a falsity-maker of the atomic sentence. To see that this corresponds to [Cut], it is helpful to read [Cut] contrapositively: if $\Gamma \not\succ \Delta$, then either $\Gamma \not\succ p, \Delta$ or $\Gamma, p \not\succ \Delta$. If we think of this in terms of TM-validity, it says that if there is a possible fusion of truth-makers of Γ and falsity-makers of Δ, then there is either a

possible fusion of truth-makers of $\Gamma \cup \{p\}$ and falsity-makers of Δ or there is a possible fusion of truth-makers of Γ and falsity-makers of $\Delta \cup \{p\}$. Exhaustivity ensures that this holds; so it ensures that [Cut] is sound.

Moreover, assuming the expressibility of propositions again, Exhaustivity is not only sufficient but also necessary for the soundness of [Cut]. For suppose that there is a possible state, u, such that neither $\exists s \in |p|^{+}$ $(u \Cup s \in S^{\diamond})$ nor $\exists t \in |p|^{-}$ $(u \Cup t \in S^{\diamond})$. If u is the only truth-maker of U and we let u be possible, then $U \not\models_{TM}$ but $U, p \models_{TM}$ and $U \models_{TM} p$. And that is a counterexample to [Cut]. Hence, Exhaustivity and [Cut] are equivalent with respect to their effect on consequence.

So, not only do the operational rules of our sequent calculus correspond exactly to the semantic clauses in truth-maker theory but the structural rules correspond exactly to constraints on possible states. We can summarize our results regarding the correspondence between structural rules and constraints on possible states in the following proposition.

Proposition 43. *The rules [Containment], [Weakening], and [Cut] preserve TM-validity if and only if possible states obey Exclusivity, Downward-Closure, and Exhaustivity respectively. (Appendix, Proposition 53)*

Putting these results together suffices to show that $\models_{LTM}$ coincides with classical propositional consequence, because the sequent calculus that includes all of our operational and structural rules is a formulation of classical propositional logic, $\models_{CL}$, and it is also sound and complete with respect to $\models_{LTM}$.

Proposition 44. *If we impose Exclusivity, Exhaustivity, and Downward-Closure, then* $\Gamma \models_{LTM} \Delta$ *if and only if* $\Gamma \models_{CL} \Delta$. *(Appendix, Proposition 54)*

Of course, our goal here is not to recover classical logic in truth-maker theory. Rather, the interest of this result lies in the fact that we have obtained it by finding a correspondence between the parts of our truth-maker theory and our sequent calculus. We now have a clearer picture of how truth-maker theory and the sequent calculus rules from the previous chapter correspond to each other. And we can now turn to the question how we can codify open reason relations in truth-maker theory.

4.4 Accommodating Open Reason Relations

We have seen in the previous section that the structural rules of [Cut] and [Weakening] correspond to constraints on possible states that are familiar from Fine, namely Exhaustivity and Downward-Closure.[34] This suggests

that if we want to allow for reason relations that are open in the sense that they are not constrained by Transitivity or Monotonicity, we can do that in truth-maker theory by dropping Exhaustivity or Downward Closure. The aim of this section is to show that this is indeed the case.

4.4.1 Nontransitive Truth-Maker Consequence

Let us then formulate a nontransitive version of truth-maker consequence by dropping Exhaustivity. If we consider logical truth-maker consequence, then this change by itself won't have any effect on our logical consequence relation because [Cut] is admissible. In terms of truth-maker theory: Exhaustivity is redundant if Exclusivity is the only constraint on possible states that holds in all models. There is, however, a well-known way to make [Cut] fail in a useful and interesting way, namely the so-called strict/tolerant logic STT, which includes a transparent truth-predicate and was developed as a response to the semantic paradoxes (Cobreros et al., 2013, 2012; Ripley, 2012). Hence, recovering STT in truth-maker theory not only gives us a new kind of semantics for nontransitive logics but is also an interesting test-case for our idea that we can allow for nontransitive reason relations in truth-maker theory by dropping Exhaustivity.[35]

In the absence of any additional notions like truth, the logic ST is equivalent to the NMMS extension of a base that consists of all and only the instances of Containment for a given atomic language, which we have seen to coincide with classical logic in the previous chapter. This logic can be extended to STT by adding a truth predicate, as follows: Let's add to our object language a canonical name $\bar{A}$ for every sentence A and a truth-predicate, Tr, for which we add the following sequent rules to NMMS.[36]

$$\frac{\Gamma, A \succ \Delta}{\Gamma, Tr(\bar{A}) \succ \Delta}\ [\mathrm{Lt}] \qquad\qquad \frac{\Gamma \succ A, \Delta}{\Gamma \succ Tr(\bar{A}), \Delta}\ [\mathrm{Rt}]$$

Definition 45 (STT-calculus)**.** The STT-calculus is the sequent calculus that is like NMMS, over a base that consists of all and only the instances of Containment for the atomic sentences, except that it includes the rules [Lt] and [Rt]. We say that $\Gamma \vdash_{\mathsf{STT}} \Delta$ if and only if the sequent $\Gamma \succ \Delta$ is derivable in this calculus.[37]

Let us also allow for self-reference by allowing sentences that include their own names. Thus, we can formulate a Liar sentence, $\neg Tr(\bar{\lambda})$, whose name is $\bar{\lambda}$. This sentence says of itself that it is not true. Note that λ is everywhere intersubstitutable with $\neg Tr(\bar{\lambda})$ *salva consequentia*.[38] Since $Tr(\bar{\lambda})$ is an atomic sentence, [Containment] yields $Tr(\bar{\lambda}) \succ Tr(\bar{\lambda})$. Using

[Lt], [Rt], our negation rules, and the intersubstitutability of $\neg Tr(\bar{\lambda})$ and λ, we can derive $\succ\lambda$ and $\lambda\succ$ (Ripley, 2013). Applying [Cut] now yields the empty sequent. However, [Cut] is not a rule of STT. Hence, we can accept that $\succ\lambda$ and $\lambda\succ$ are derivable without this rendering our consequence relation trivial. This is the essence of the nontransitive approach to the Liar Paradox.

If the isomorphism that we explained above lives up to our promises above, we can now make the analogous changes in our truth-maker theory. And, indeed, we can do this (see Hlobil, 2022a).[39] Namely, we can do this by adding a truth predicate to our truth-maker theory, which is governed by the following semantic clauses:

(tr+) $s \Vdash Tr(\bar{A})$ if and only if $s \Vdash A$

(tr−) $s \dashv\!\Vdash Tr(\bar{A})$ if and only if $s \dashv\!\Vdash A$

We allow for self-reference as in STT, so that the language of our truth-maker theory includes a liar sentence $\neg Tr(\bar{\lambda})$, whose name is $\bar{\lambda}$. If our truth-maker theory includes Exhaustivity, this implies that all states in all models are impossible. For, suppose that $s \Vdash \lambda$, that is, $s \Vdash \neg Tr(\bar{\lambda})$. Then $s \dashv\!\Vdash Tr(\bar{\lambda})$ and, hence, $s \dashv\!\Vdash \lambda$. So, a state makes the liar sentence true if and only if it makes the liar sentence false. Thus, by Exclusivity, any state that includes a truth-maker or a falsity-maker of the liar sentence is impossible. But if we can extend every possible state to a possible state by either adding a truth-maker or a falsity-maker of the liar sentence but there is no possible state that includes a truth-maker or a falsity-maker of the liar sentence, then there cannot be any possible state. Given the definition of truth-maker consequence, it follows that every set of premises implies every set of conclusions. The solution to this paradox that mirrors the solution of STT is, of course, to reject Exhaustivity. Let us define this version of truth-maker theory explicitly.

Definition 46 (TM-consequence for base STT). Let $\overset{STT}{\underset{LTM}{\models}}$ be the consequence relation just like $\underset{LTM}{\models}$ except that we include models that violate Exhaustivity and the language has a truth-predicate whose interpretation obeys the clauses (tr+) and (tr−).

These changes have exactly the effects that we have envisaged. The semantic clauses have the same effect on consequence as the sequent rules for the truth-predicate, and dropping [Cut] and Exhaustivity have also the same effect. Hence, this version of (logical) truth-maker consequence coincides with the consequence relation of STT.

Proposition 47. $\Gamma \vdash_{\mathsf{STT}} \Delta$ *if and only if* $\Gamma \overset{STT}{\underset{LTM}{\models}} \Delta$. *(Appendix, Proposition 59)*

Notice that the two formal systems don't only match in their consequence relations. Rather, they match in a piece-by-piece fashion, in the way explained in the previous subsection. Indeed, all we did was to use the correspondence from the previous subsection to formulate a version of truth-maker consequence that mirrors the sequent calculus for STT.[40]

We can illustrate this by spelling out how the philosophical interpretation of failures of [Cut] can be translated from normative bilateralism into truth-maker theory. For normative bilateralism, the intuitive idea behind the rejections of [Cut] is that any position that includes an assertion or a denial of the liar sentence is out-of-bounds. When we translate this into truth-maker bilateralism, the result is this: Any state that includes a verifier or a falsifier of the liar sentence is impossible. Given any possible state, this yields a violation of Exhaustivity. Since the world is a possible state, we can express the idea by saying that the world cannot contain anything that makes the liar sentence either true or false. Just as we should neither assert nor deny the liar sentence, so the world can neither verify nor falsify it.

We have thus formulated a truth-maker semantics for strict/tolerant logic, and this illustrates how we can allow for nontransitive reason relations in truth-maker theory. Now, the reason why we rejected transitivity in the previous chapters had nothing to do with the semantic paradoxes. We argued rather that once we free ourselves from the constraint of Monotonicity on reason relations, our language will have crippling expressive limitations unless we also allow for nontransitive reason relations. Hence, we must consider whether the strategy for allowing for open reason relations in truth-maker theory can yield the desired results when we move away from STT as a test-case.

4.4.2 Nonmonotonic Truth-Maker Consequence

We now want to make the final step to show that the framework of truth-maker consequence is isomorphic to the theory that we presented in the previous chapter. We want to show that if we reject Monotonicity and Transitivity in our truth-maker theory and we can add material implications among atomic sentences, then the result is the same as the logical extension of these material implications by the rules of NMMS from the previous chapter.

As we have seen in the previous chapter, NMMS allows for failures of Monotonicity and Transitivity. We have seen in the previous section that Monotonicity corresponds to Downward-Closure in truth-maker theory, and Transitivity corresponds to Exhaustivity. Hence, in order to

mirror NMMS within truth-maker theory, we drop Downward-Closure and Exhaustivity. Moreover, we must add material, non-logical implications among atomic sentences to our truth-maker theory in order to mirror the use of base vocabularies in NMMS.

Recall that we started in the previous chapter with a material base vocabulary, which consists of an atomic language $\mathfrak{L}_{\mathfrak{B}}$ and a consequence relation $\mid\!\sim_{\mathfrak{B}}$ over subsets of that atomic language. The base consequence relation encodes non-logical reason relations of both fundamental varieties: reasons-for and reasons-against. Let us quickly rehearse some examples of the ideas from the previous chapter. Let's suppose that c = "This is a chair" implies s = "You can sit on this." Hence, $c \mid\!\sim_{\mathfrak{B}} s$. Since this implication is defeated if we add the additional premise b = "This is broken," the corresponding pair is not in the material base. That is, $c, b \not\mid\!\sim_{\mathfrak{B}} s$. Similarly, since "This is a chair" is incompatible with v = "This is a violin," we have $c, v \mid\!\sim_{\mathfrak{B}}$. But since this incoherence is cured by adding a = "This is part of an art project that makes pieces of furniture that are musical instruments," we say $c, v, a \not\mid\!\sim_{\mathfrak{B}}$. Thus, our two reason relations, implication and incompatibility, are both defeasible, and we can model this in our formalism with an appropriate material base consequence relation. This is all familiar from the previous chapters. Moreover, we showed in the previous chapter how we can add logical vocabulary to such a material base, namely by closing the material base under the operational rules of NMMS. We saw in the previous chapter that if we restrict our material bases to those that obey Containment, all consequence relations defined by such a calculus include all of classical logic. And we showed how to add various kinds of object language operators to make explicit, in the object language, various local features of such consequence relations.

Our question now is how we can find truth-maker formulations of these consequence relations. The first step is to ensure that a truth-maker consequence relation includes all the atomic sequents in a given material base $\mathfrak{B}$. We can do that by stipulating, for every pair $\langle \Gamma, \Delta \rangle \in \mid\!\sim_{\mathfrak{B}}$, that any fusion of verifiers for everything in Γ and falsifiers for everything in Δ is impossible. The following definition does that.

Definition 48 (TM-consequence for base $\mathfrak{B}$). $\overset{\mathfrak{B}}{\models_{TM}}$ is the consequence relation that is just like $\models_{TM}$ except that the only constraint on possible states is that $s \notin S^{\diamond}$ if and only if $s = \Cup\{g_1, ..., g_n, d_1, ..., d_m\}$ such that $\langle \Gamma, \Delta \rangle \in \mid\!\sim_{\mathfrak{B}}$ and $\Gamma = \{\gamma_1, ..., \gamma_n\}$ and $\Delta = \{\delta_1, ..., \delta_m\}$ and $\forall 1 \leq i \leq n\, (g_i \in |\gamma_i|^{+})$ and $\forall 1 \leq i \leq m\, (d_i \in |\delta_i|^{-})$.

This definition says that, for every sequent in the material base, every state that the sequent deems impossible is impossible in our model. We

could quantify over all truth-maker models that meet this constraint, but we don't really need to do that in order to define $\frac{\mathfrak{B}}{\models_{TM}}$. The model in which all and only those states are impossible that must be impossible for the material base to hold already includes all the counterexamples to all invalid sequents.

If we understand consequence in this way, then the failures of Monotonicity in the examples above look as follows: In the absence of any defeating facts, the state of something being a chair and the state of one being unable to sit on it are incompatible, that is, their fusion is impossible. However, if we add to this fusion the state that the object in question is broken, then the states are no longer incompatible, that is, the fusion of all three states is a possible state. Similarly, the states of something being a chair and the state of it being a violin are incompatible, that is, their fusion is an impossible state. However, if we fuse this state with the state of the object being part of an art project that makes pieces of furniture that are musical instruments, then the resulting state isn't impossible. The thought is that certain states can make states compatible with each other that are otherwise incompatible. These states need, as it were, the further state in order to fit together.

We know that the semantic clauses in truth-maker theory correspond to the operational rules in sequent calculi. So adding logical vocabulary that obeys the semantic clauses has the same effect on consequence as closing the base sequents under the corresponding sequent rules. And since we didn't enforce Exhaustivity and Downward-Closure but insisted on Exclusivity, the resulting consequence relation is guaranteed to obey Containment, but it may be nontransitive and nonmonotonic. In fact, the result is exactly what you might expect: the sequent calculus version of our current logic and the truth-maker consequence relation of it coincide if the material bases are the same.

Proposition 49. $\Gamma \succ \Delta$ *is derivable in* $\mathsf{NMMS}_{\mathfrak{B}}$ *just in case* $\Gamma \frac{\mathfrak{B}}{\models_{TM}} \Delta$. *(Appendix, Proposition 60)*

So we now have a way to construct a truth-maker semantics for any nonmonotonic logic that can be obtained by closing an arbitrary set of atomic sequents under given operational rules. Thus, we can formulate our account of consequence relation in terms of the rules of NMMS from the previous chapter within truth-maker theory in a surprisingly straightforward way. We simply take the semantic clauses that correspond to the desired sequent rules, and we stipulate the material base by way of a constraint on possible states. If we restrict ourselves to bases that obey Containment, then the result will include all classically valid sequents. If we want to enforce a structural rule, we add the corresponding constraint

on possible states. If we want to add another bit of logical vocabulary, we read off the semantic clauses from the sequent rules. The clause for truth-makers is given by the left rule, and the clause for falsity-makers is given by the right rule. We saw an example of this when we looked at the clauses for the truth-predicate.

The isomorphism between the theories ensures that parallel tweaks yield parallel results. In this way, we can see which semantic-representationalist rules and principles correspond to the rules and principles of a particular pragmatic-normative account of content and reason relations, and *vice versa*. And what connects the two sides of this correspondence to each other is the notion of representation. For, a particular pragmatic-normative account corresponds to a particular semantic-representationalist account just in case the rational form attributed to a sentence by the pragmatic-normative account is the same as the rational form attributed to the worldly proposition expressed by that sentence according to the semantic-representationalist account.

The upshot of all this is that the sentences that occur in discursive acts that are governed by the norms that are codified in a consequence relations defined by $\mathsf{NMMS}_{\mathfrak{B}}$ share their rational form with the worldly propositions that these same sentences express according to the truth-maker model that defines the corresponding relation $\overset{\mathfrak{B}}{\underset{TM}{\models}}$. That is, the modal relations of exclusion among assertions and denials of sentences and among their truth-makers and falsity-makers is isomorphic, supposing that we use the same base relation in $\mathsf{NMMS}_{\mathfrak{B}}$ and in truth-maker theory. Thus, we can now see how a pragmatic-normative account of content and reason relations can be isomorphic to a semantic-representationalist account of content and reason relations. And we can think of the relation of representation as holding precisely if and when such an isomorphism indeed holds in a particular case. Moreover, we have seen how we can ensure that such an isomorphism holds for logically complex sentences, supposing that it holds for atomic sentences. All this can be done no matter how radically open-structured (substructural) the reason relations that articulate the base vocabulary are.

4.5 Conclusion

We began the Introduction to this book by pointing out that two principal traditions in the philosophy of language are distinguished by whether they focus to begin with on the *use* of linguistic expressions, or on their *meanings*. These topics are addressed by two different disciplines: pragmatics (in a broad sense) and semantics. We urged thinking about these in terms of the metavocabularies they use to discuss the aspects of the base vocabularies they target. In Chapter One, we presented our version of a

bilateral normative pragmatic metavocabulary. It specifies norms governing discursive practices of using declarative sentences in speech acts of asserting and denying, which manifest doxastic attitudes of acceptance and rejection. In the present chapter, we discussed the sophisticated contemporary modal truth-maker formal semantic metavocabulary. It appeals both to matter-of-factual worldly states, divided into the possible and the impossible, and to their mereological relations to one another, to determine what makes declarative sentences true or false, depending on how things objectively are. Though the pragmatic and semantic theoretical enterprises can be, and have been, conducted each in their own terms, substantially independently of each other, it is also important to understand the relations between the activities of talking and thinking (using expressions as representings of something) and what is being talked or thought about (what is thereby represented). Thought of in these terms, understanding the relations between what is expressed by pragmatic metavocabularies and what is expressed by semantic metavocabularies is essential to understanding the intentional nexus between (attempted) knowings and what is known.

We have adopted an indirect approach to this challenge. We regard (the right kind of) pragmatic and semantic metavocabularies as providing two different conceptual and explanatory perspectives on a common object: reason relations of implication and incompatibility. They should be understood as being metavocabularies specifying the reason relations of some autonomous base vocabulary. They can therefore be thought of as metavocabularies of reason, or rational metavocabularies, in the sense in which we use that term. On the pragmatic side, the two kinds of reason relation determine what claims should be accorded the practical significance of being reasons for and against what others (in the broad sense of "reasons" operative in this book). They articulate (aspects of) standards for assessing the success of speech acts of justifying, seeking rationally to secure entitlements to constellations of commitments to accept or reject, and of speech acts of challenging, rationally contesting such entitlements. On the semantic side, the two kinds of reason relation show up as modal relations determining how worldly states are possible or impossible relative to one another.

Of course each sort of metavocabulary also talks about things that are not reason relations: about speech acts and their relations to two kinds of normative status, in the one case, and about states and their fusions and their absolute and relative possibility and impossibility, in the other case. It is that further conceptual material (extrinsic to the base or object-language vocabulary for which they are metavocabularies) that pragmatic and semantic theories in these metavocabularies deploy to *explain* relations of implication and incompatibility, each in its own way. Each offers an account of what relations of rational implication and incompatibility are,

and of what determines what implies and is incompatible with what. The pragmatic and semantic explanations of reason relations differ not only in the conceptual raw materials they appeal to, but also in the kind of modal relations that are essential to the distinct sorts of account they articulate. The fundamental modal vocabulary employed by the pragmatics consists of normative terms such as "commitment" and "entitlement." The fundamental modal vocabulary employed to specify the metaphysical universe in terms of which semantic interpretants are defined consists of alethic terms such as "possible" and "impossible."

In this chapter we showed that in spite of these substantial differences in explanatory raw materials and form there is a structural correspondence between the pragmatics-first approach as we pursued it in the previous two chapters, and a semantics-first approach as it is pursued in truth-maker theory. Furthermore, the correspondence is a conceptually and technically fruitful one. On the pragmatic side, we fill in the discursive social practices that confer content on the notion of "out-of-boundness," and further articulate the distinction and relations between the two kinds of normative status implicit in bilateralism: the commitments that make up a position and the notion of normative impossibility or "out-of-boundness" understood as preclusion of entitlement to such a constellation of commitments. By offering a truth-maker semantics for the expressively LX logic NMMS, we show how the truth-maker framework can be extended in a philosophically motivatable and formally tractable way to codify radically substructural nonmonotonic and nontransitive reason relations of implication and incompatibility. When so used, the apparatus tells us exactly what the metaphysical structure of worldly states must be like in order properly to be thought about using nonmonotonic and nontransitive reasoning practices.

In fact, what is on offer is more than just a correspondence. We have shown how to define a detailed isomorphism relating the pragmatic and semantic metavocabularies on which we have focused. They are isomorphic precisely at the level of reason relations. That isomorphism defines higher-order representational norms governing the relations between the pragmatic use of declarative sentences and their semantic meanings. Very roughly, in order to represent worldly propositions with sentences, the conceptual norms governing discursive practice must be such that according to them, the pragmatic-normative structure *ought* to match the alethic-modal structure. More specifically, properly to represent relations of worldly consequence, language users should in practice take it that they cannot be entitled to accept all the sentences in a set Γ and reject all the sentences in Δ just in case the fusion of any set of truth-makers of all of Γ with any set of false-makers of all of Δ is always an impossible state. (On the side of incompatibility, sets of commitments to accept to which one

cannot be jointly entitled should line up with sets of sentences such that every fusion of truth-makers for each of them is an impossible state.)

We have been using the term "rational form" to talk about the roles sentences and what the sentences express play with respect to reason relations. Insofar as the higher-order representational norm defined by that isomorphism holds across the board, the rational form conferred on sentences by their use as specified in the normative bilateral metavocabulary and the rational form conferred on sentences by their worldly truth-makers and false-makers coincide. That is what makes it possible for declarative sentences to say of what is, that it is, and of what is not, that it is not. That discursive acts and worldly states can share their rational forms allows us to see that our discursive acts (representings) and the worldly states that they are about (represent) can have something in common. And what they have in common is that in virtue of which our discursive acts have the contents they have and in virtue of which the worldly states are the kinds of states they are. The rational pragmatic-semantic isomorphism accordingly fills in important details of what in the Introduction we called "bimodal conceptual realism."

The vantage point afforded by defining this isomorphism shows how to understand as complementary rather than competing accounts the very different explanations pragmatic and semantic metavocabularies offer of what it is for sets of sentences to imply or be incompatible with other sets of sentences. For it permits the precise specification of the topic common to them, the object on which they offer complementary perspectives, what they turn out both to present, albeit in different guises—namely reason relations and the rational forms they functionally define. Isomorphisms define equivalence relations that support the introduction by abstraction of new terms corresponding to equivalence classes of expressions in the vocabularies from which they are abstracted. Though we found our way to the conception of reason relations via the different specifications of it in (the right kind of) pragmatic and semantic metavocabularies, we are now in a position to move beyond those perspectives on reason relations. Reason relations of implication and incompatibility, and the rational forms they define and articulate, are introduced by abstraction, picked out by equivalence classes with respect to the pragmatic-semantic isomorphism. They are what has been our topic all along. The constructions and demonstrations of this chapter have advanced the story by putting us in a position properly to delineate the objects these different metavocabularies provide perspectives on, in a way that is unalloyed by admixture with extraneous pragmatic or semantic concepts, important and informative as those otherwise are.

The abstract notion of rational form brought into view by the isomorphism between consequence relations specified in deontic normative

pragmatic and alethic modal semantic metavocabularies suggests that we should investigate those abstract common forms in their own right. For if what we have said so far is correct, this common form is the structure of reason relations in general—not only as they appear in our discursive practice of giving and asking for reasons but also as they articulate the worldly relations about states to which our conceptual norms correspond. Accordingly, in the next chapter we introduce a metavocabulary that is an intrinsic metavocabulary specifying rational forms understood as roles with respect to reason relations considered abstractly: conceptual propositional roles. We provide an implication-space semantics for conceptual roles. It is a way of working out the aspirations of semantic inferentialism, as the construction of NMMS as universally LX (extending even to arbitrary open-structured base vocabularies) in Chapter Three is a way of working out the aspirations of logical expressivism.

In calling the implication-space model-theoretic semantics for reason relations and conceptual roles an "intrinsic" metavocabulary, we are contrasting it with both pragmatic and semantic metavocabularies. They count as "extrinsic" metavocabularies because their explanations of reason relations appeal to conceptual raw materials that are not drawn entirely from the base vocabularies whose reason relations are the targets of their explanations. In Chapter Three we considered logical vocabulary as being elaborated entirely from those base vocabularies, importing no extrinsic material to them. In using the base vocabularies practitioners already know how to do everything they need to know how to do to deploy the logical vocabulary. Logical vocabulary is in this sense an intrinsic metavocabulary for expressing the reason relations of base vocabularies it extends. We saw that its distinctive function is expressive and explicative rather than explanatory. Logical vocabulary serves to make explicit the reason relations of the base vocabulary, in the sense of saying what follows from what and what is incompatible with what, rather than aiming at explaining why or in what sense things stand in the reason relations they do—for instance, by appeal to features of discursive practice, or of worldly states. The implication-space semantics characterizes rational forms (propositional contents) directly in terms of roles with respect to reason relations. It, too, is an intrinsic-explicative rational metavocabulary, rather than an extrinsic-explanatory one. It completes the set of four rational metavocabularies we employ in this book as tools for specifying, investigating, and manipulating reason relations and conceptual roles.

Our discussion in the next chapter of the implication-space semantics for conceptual roles also adds two further connections among those four kinds of rational metavocabulary. In this chapter we specified two such relations precisely, in the form of the isomorphism between pragmatic and semantic metavocabularies, and the construction of a truth-maker

semantics that is sound and complete for the logical vocabulary of NMMS. Chapter Five shows how to define a more abstract implication-space version of any truth-maker semantic metavocabulary, and also shows that it, too, is sound and complete for the logic NMMS, even for radically open-structured nonmonotonic and hypernonmonotonic, nontransitive base vocabularies. The chapter closes by building on the implication-space semantic vocabulary a still more abstract characterization of conceptual roles, including partial conceptual roles such as premissory and conclusory roles. There we show how to use that metavocabulary to draw connections between this abstract form of reason relations and a number of familiar and interesting logics, including Linear Logic, Priest's Logic of Paradox, and Strong Kleene Logic.

The four metavocabularies for specifying reason relations and conceptual roles or rational forms are divided into the two extrinsic-explanatory rational metavocabularies, bilateral-pragmatic and truth-maker-semantic, and the two intrinsic-explicative rational metavocabularies, logical and implication-space. Along with the criss-crossing relations among them that we demonstrate and explore, this is the structure that we propose for a functionalist specification of reason relations and rational forms, in terms of the role those phenomena play as the objects of all these intricately interconnected kinds of metavocabulary. Proceeding in this way is pursuing the strategy that above we called "meta-rational functionalism": defining the object of study by the reason relations that articulate the meta-meta-vocabulary (metavocabulary for rational metavocabularies) that is the use-language of this book. We turn now to consideration of the abstract implication-space semantics for conceptual roles that is the fourth vertex of the tetrahedral structure of rational metavocabularies.

4.6 Appendix

Lemma 50. *For every application of an operational rule of* NMMS, *the set of states deemed impossible by the bottom-sequent is the union of the states deemed impossible by the top-sequents.*

Proof. We do the case for conjunction; the proofs for negation and disjunction are analogous. For [L$\wedge$]: Note that by our semantic clauses $|A \wedge B|^+ = \{s : \exists a \in |A|^+ \exists b \in |B|^+ (s = a \Cup b)\}$. Hence, for any Γ and Δ, we have $\{g \Cup d \Cup a \Cup b : g \in |\bigwedge \Gamma|^+ \text{ and } d \in |\bigvee \Delta|^- \text{ and } a \in |A|^+ \text{ and } b \in |B|^+\} = \{g \Cup d \Cup s : g \in |\bigwedge \Gamma|^+ \text{ and } d \in |\bigvee \Delta|^- \text{ and } s \in |A \wedge B|^+\}$. So the states deemed impossible by $\Gamma, A, B \succ \Delta$ are identical to those deemed impossible by $\Gamma, A \wedge B \succ \Delta$.

Similarly for [R$\wedge$], note that $|A \wedge B|^- = |A|^- \cup |B|^- \cup \{s : \exists a \in |A|^- \exists b \in |B|^- (s = a \Cup b)\}$. Therefore, $\{g \Cup d \Cup s : g \in |\bigwedge \Gamma|^+ \text{ and}$

$d \in |\bigvee \Delta|^-$ and $s \in |A \wedge B|^-\} = \{g \Cup d \Cup a : g \in |\bigwedge \Gamma|^+$ and $d \in |\bigvee \Delta|^-$ and $a \in |A|^-\} \cup \{g \Cup d \Cup b : g \in |\bigwedge \Gamma|^+$ and $d \in |\bigvee \Delta|^-$ and $b \in |B|^-\}$ $\cup \{g \Cup d \Cup c : g \in |\bigwedge \Gamma|^+$ and $d \in |\bigvee \Delta|^-$ and $c \in \{x : \exists a \in |A|^- \exists b \in |B|^- (x = a \Cup b)\}\}$. Hence, the states deemed impossible by $\Gamma \succ \Delta, A \wedge B$ are the union of those deemed impossible by $\Gamma \succ \Delta, A$ and $\Gamma \succ \Delta, B$ and $\Gamma \succ \Delta, A, B$. ■

Proposition 51. *All operational rules of* NMMS *are sound for* $\models_{TM}$*, that is, if all the top-sequents are TM-valid, then so is the bottom sequent.*

Proof. Immediate from Lemma 50. ■

Proposition 52. *If a model is a counterexample to the top-sequent of an application of an operational rule of* NMMS*, then the model is also a counterexample to the bottom-sequent.*

Proof. By Lemma 50, if a state deemed impossible by a top-sequent is possible in $\mathcal{M}$, then a state deemed impossible by the bottom sequent is possible in $\mathcal{M}$. ■

Proposition 53. *The rules [Containment], [Weakening], and [Cut] are sound for* $\models_{TM}$ *if and only if possible states obey Exclusivity, Downward-Closure, and Exhaustivity respectively.*

Proof. Downward-Closure and [Weakening]: Suppose that $\Gamma \models_{TM} \Delta$. Then any fusion of verifiers of everything in Γ and falsifiers of everything in Δ is impossible. By Downward-Closure, all states that include any such fusion as a part are also impossible. Hence, $\Theta, \Gamma \models_{TM} \Delta, \Sigma$. For the other direction, suppose that [Weakening] is sound and that $s \in S^\diamond$ and $t \sqsubseteq s$. In accordance with Assumption 39, let the sentence S be such that s is the unique truth-maker of S, and let the sentence T be such that t is the unique truth-maker of T. Since $s \in S^\diamond$ and $t \sqsubseteq s$, we know that $s \Cup t = s$ and $S, T \not\models_{TM}$. By the contrapositive of [Weakening], $T \not\models_{TM}$. Hence, $t \in S^\diamond$.

Exclusivity and [Containment]: Let $s \in |p|^+$ and $t \in |p|^-$, then Exclusivity implies $\Gamma, p \models_{TM} p, \Delta$. So [Containment] is valid. For the other direction, by the validity of [Containment], for any Γ and Δ, we have $\Gamma, p \models_{TM} p, \Delta$. So $\forall u(s \Cup t \Cup u \notin S^\diamond)$.

Exhaustivity and [Cut]: Suppose that $\Gamma \not\models_{TM} \Delta$ and let u be a state witnessing this fact, i.e., a state that is a fusion of verifiers of everything in Γ and falsifiers of everything in Δ such that $u \in S^\diamond$. By Exhaustivity, $\exists s \in |p|^+ (u \Cup s \in S^\diamond)$ or $\exists t \in |p|^- (u \Cup t \in S^\diamond)$. Therefore, either $p, \Gamma \not\models_{TM} \Delta$ or $\Gamma \not\models_{TM} \Delta, p$. But that is just what is required for the

contrapositive of [Cut]. For the other direction, suppose that [Cut] is valid for $\models_{TM}$. Let u be possible; and, in accordance with Assumption 39, let Γ be a set such that u is the unique state that is a fusion of verifiers for everything in Γ. Hence, $\Gamma \not\models_{TM}$. By the validity of [Cut], either $p, \Gamma \not\models_{TM}$ or $\Gamma \not\models_{TM} p$. Hence, either $\exists s \in |p|^{+}\ (u \Cup s \in S^{\diamond})$ or $\exists t \in |p|^{-}\ (u \Cup t \in S^{\diamond})$. ■

Proposition 54. *If we impose Exclusivity, Exhaustivity, and Downward-Closure, then* $\Gamma \models_{LTM} \Delta$ *if and only if* $\Gamma \models_{CL} \Delta$.

Proof. It suffices to show that NMMS over a base that consists of all and only the atomic instances of Containment is sound and complete with respect to both consequence relations, $\models_{LTM}$ and $\models_{CL}$. We have seen in the previous chapter that NMMS over a base that consists of all and only the atomic instances of Containment is sound and complete with respect to $\models_{CL}$. For $\models_{LTM}$, we know that soundness holds from Propositions 53 and 51. For completeness, suppose that there is no proof of $\Gamma \succ \Delta$. Hence, a proof-search for $\Gamma \succ \Delta$ yields an atomic sequent, $\Gamma_0 \succ \Delta_0$, where $\Gamma_0 \cap \Delta_0 = \emptyset$. Let $\mathcal{M}$ be a model in which $s \in S^{\diamond}$ and $s = u \sqcup t$ such that $u \Vdash \Gamma$ and $t \dashv\!\Vdash \Delta$. This is a counterexample to $\Gamma_0 \succ \Delta_0$. By Proposition 52, it follows that it is also a counterexample to $\Gamma \succ \Delta$. Since $\Gamma_0 \cap \Delta_0 = \emptyset$, such a model isn't ruled out by Exclusivity, which is the only principle that could rule out such a model. So, $\mathcal{M}$ shows that $\Gamma \not\models_{TM} \Delta$. ■

The relation between STT and the truth-maker theory in this chapter is investigated in (Hlobil, 2022a). Here we merely sketch the proof-ideas. In order to prove the completeness of the STT calculus with respect to $\models^{STT}_{LTM}$, we cannot use the technique of proof-searches because proof-searches are no longer guaranteed to terminate. Hence, we follow Ripley (2013, 162–63) in using the technique of (possibly infinite) reduction trees from Takeuti (1987).

Definition 55 (Reduction tree)**.** The reduction tree for a sequent, $\Gamma \succ \Delta$, is the possibly infinite tree that results from starting with $\Gamma \succ \Delta$ as the root of the tree and then extending at each stage each top-most sequent of the tree as follows, until all branches are closed or else extending the tree ω-many times: (i) If the sequent is an axiom, i.e., is such that the left and the right side share an atomic sentence, then the branch remains unchanged and is closed. (ii) If the sequent has the form $\Gamma, \neg A \succ \Delta$ or $\Gamma \succ \neg A, \Delta$ and no reduction has been applied to $\neg A$ in previous stages, they reduce to $\Gamma, \neg A \succ A, \Delta$ and $\Gamma, A \succ \neg A, \Delta$, respectively. (iii) If the sequent has the form $\Gamma, A \wedge B \succ \Delta$ and no reduction has been applied to $A \wedge B$ in previous stages, it reduces to $\Gamma, A, B, A \wedge B \succ \Delta$; and if it has the form $\Gamma \succ A \wedge B, \Delta$ and no reduction has been applied to $A \wedge B$ in previous stages, the reduction tree

branches into $\Gamma \succ A \wedge B, A, \Delta$ and $\Gamma \succ A \wedge B, B, \Delta$ and $\Gamma \succ A \wedge B, A, B, \Delta$. (iv) Similarly, $\Gamma \succ A \vee B, \Delta$ reduces to $\Gamma \succ A \vee B, A, B, \Delta$; and $\Gamma, A \vee B \succ \Delta$ reduces to $\Gamma, A, A \vee B \succ \Delta$ and $\Gamma, B, A \vee B \succ \Delta$ and $\Gamma, A, B, A \vee B \succ \Delta$. (v) $\Gamma, Tr(\bar{A}) \succ \Delta$ reduces to $\Gamma, A, Tr(\bar{A}) \succ \Delta$; and $\Gamma \succ Tr(\bar{A}), \Delta$ reduces to $\Gamma \succ Tr(\bar{A}), A, \Delta$.

Lemma 56. *The set of states deemed impossible by a sequent is the union of the states deemed impossible by the sequents to which it reduces in a reduction tree.*

Proof. We look at each clause in the reduction procedure. The lemma holds trivially for clause (i). It holds for (ii) because the truth-makers of $\neg A$ are exactly the falsity-makers of A, and vice versa. The other cases, in particular those for (v), are analogous except for when the reduction tree branches out, such as in the case of $\Gamma \succ A \wedge B, \Delta$. In this case, the lemma holds because the falsity-makers of $A \wedge B$ are the union of the falsity-makers of A, the falsity-makers of B, and any fusion of such falsity-makers, which corresponds to the three sequents that result from the reduction. ■

Definition 57 (Sequents resulting from an open branch of a reduction tree)**.** If an open branch of a reduction tree terminates, the resulting sequent is the leaf of that branch. If the open branch does not terminate, then the resulting sequent is the sequent $\Gamma_\omega \succ \Delta_\omega$, where Γ_ω is the union of all the sets on the left side of sequents in this open branch and and Δ_ω is the union of the sets on the right side of sequents in the branch.

Lemma 58. *Let $\Gamma \succ \Delta$ be a sequent resulting from an open branch, let Γ^{at} be the set of atomic sentences in Γ, and let Δ^{at} be the set of atomic sentences in Δ. Then a state that is deemed impossible by $\Gamma^{at} \succ \Delta^{at}$ includes as a part a truth-maker for every sentence in Γ and a falsity-maker for every sentence in Δ.*

Proof. We argue by induction on the complexity of sentences in $\Gamma \cup \Delta$. The states deemed impossible by $\Gamma^{at} \succ \Delta^{at}$ trivially include truth-makers for every atomic sentence in Γ and falsity-makers for every atomic sentence in Δ. Suppose our lemma holds for sentences up to complexity n, and let us consider sentences of complexity $n + 1$. Note that since we have an open branch, we know that all possible reduction procedures have been applied. For negations in Γ, we know that the negatum, which is of complexity n, is in Δ. So, by our induction hypothesis states deemed impossible by $\Gamma^{at} \succ \Delta^{at}$ include a falsity-maker for the negatum, which is a truth-maker for our negation. Similarly for all other connectives where the reduction tree does not branch. For disjunctions on the left, we know that Γ contains also one or both of the disjunctions, which are of complexity n. So by

our hypothesis, $\Gamma^{at} \succ \Delta^{at}$ contains truth-makers for one or both disjuncts, and any of these options ensures that it includes a truth-maker for the disjunction. Similarly for conjunctions on the right. ■

Proposition 59. $\Gamma \vdash_{\mathsf{STT}} \Delta$ *if and only if* $\Gamma \models^{STT}_{LTM} \Delta$.

Proof. Left-to-right: The proofs for Proposition 53 and Proposition 51 still go through, because clauses (tr+) and (tr−) ensure that [Lt] and [Rt] are valid for $\models^{STT}_{LTM}$.

Right-to-left: Suppose that there is no proof of $\Gamma \succ \Delta$. Hence, a reduction tree for $\Gamma \succ \Delta$ has an open branch. Let $\Gamma_\omega \succ \Delta_\omega$ be the sequent that results from that branch, and let $\Gamma^{at}_\omega \succ \Delta^{at}_\omega$ be the sequent that results from $\Gamma_\omega \succ \Delta_\omega$ by omitting all complex sentences. We can use as our desired countermodel any model that makes possible one of the states deemed impossible by $\Gamma^{at}_\omega \succ \Delta^{at}_\omega$. For by Lemma 58, any state that is deemed impossible by $\Gamma^{at}_\omega \succ \Delta^{at}_\omega$ is also deemed impossible by $\Gamma_\omega \succ \Delta_\omega$. And by Lemma 56, any state deemed impossible by $\Gamma_\omega \succ \Delta_\omega$ is deemed impossible by $\Gamma \succ \Delta$. So our model is a countermodel to $\Gamma \succ \Delta$. We know that there is such a model because any model will work that makes only those states impossible that are required to be impossible by Exclusivity, and $\Gamma^{at}_\omega \cap \Delta^{at}_\omega = \emptyset$. ■

Proposition 60. $\Gamma \succ \Delta$ *is derivable in* $\mathsf{NMMS}_{\mathfrak{B}}$ *if and only if* $\Gamma \models^{\mathfrak{B}}_{TM} \Delta$.

Proof. Left-to-right: By induction on proof-height. The axioms of $\mathsf{NMMS}_{\mathfrak{B}}$ are all the sequents in $\mathrel{|\!\sim}_{\mathfrak{B}}$. Hence, they are valid according to $\models^{\mathfrak{B}}_{TM}$ by Definition 48. Our sequent rules preserve validity in TM-models by Proposition 51.

Right-to-left: Suppose that $\Gamma \succ \Delta$ is not derivable in $\mathsf{NMMS}_{\mathfrak{B}}$. A root-first proof-search must yield an atomic sequent that is not in $\mathrel{|\!\sim}_{\mathfrak{B}}$. Hence, the states that this atomic sequent deems impossible are not all impossible in all models. So we can find a model in which one of these states is possible. By Proposition 52, this is also a counterexample to $\Gamma \models^{\mathfrak{B}}_{TM} \Delta$. ■

Notes

1 In the following chapters, we will look at contents in their own right, which will ultimately also allow us to think of reason relations among contents in abstraction from relations among sentences. But in this chapter, we take sets of sentences as the relata of reason relations.

2 The ideas in this chapter have been developed in conjunction with the ideas and results presented in (Hlobil, 2022a). That paper does not, however, include any treatment of nonmonotonic consequence relations. For some of the technical

results below, it might be helpful (though not necessary) to consult (Hlobil, 2022a) and (Hlobil, 2022b).

3 We think of the contentfulness of contentful items not as consisting in a relation between these items and other items that are their contents but rather as a property of contentful items. We talk about (worldly) propositions below, but those are not contents but contentful items. So, our notion of contentful item here includes what we mean by "proposition," as will become clearer below.

4 As in the previous chapters, we restrict ourselves to declarative sentences. So we will not consider whether and how questions or imperatives might be said to represent anything.

5 This formulation is, of course, inspired by Aristotle's explanation of truth and falsity: "To say of what is that it is not, or of what is not that it is, is false, while to say of what is that it is, and of what is not that it is not, is true" (Aristotle, Meta. Γ, 1011b25).

6 In contrast to some advocates of truth-maker theory, we avoid talk about states existing or not existing to mean something like that the state obtains or does not obtain. We think that such talk invites philosophical confusion whose resolution is beyond the scope of this book.

7 This distinction is analogous to the distinction between the actual truth-value of a sentence and its truth-value in each possible world, which is often called its "intension." However, truth-maker theory does not limit us to possible worlds, and it allows us to single out the parts of reality that are responsible for the truth or falsity of a given sentence.

8 Fine (2017a, 629) uses the term "bilateral proposition" for what we here call a "worldly proposition."

9 We are here again developing an Aristotelian idea. According to Aristotle, the intellect and what it thinks about share their form (as do sensation and the sensed object).

> Within the soul the faculties of knowledge and sensation are potentially these objects, the one what is knowable, the other what is sensible. They must be either the things themselves or their forms. The former alternative is of course impossible: it is not the stone which is present in the soul but its form. It follows that the soul is analogous to the hand; for as the hand is a tool of tools, so thought is the form of forms and sense the form of sensible things. (Aristotle, De An. 431b25-432a3)

In knowing an object, the faculty of knowledge takes on the form of the known object and, hence, becomes isomorphic to that object. The nature of the intellect is that it can take on forms in this way. That is why it is the form of forms. In the special case in which what is thought about does not have any matter the known object is actually in the intellect: "For in the case of objects which involve no matter, what thinks and what is thought are identical; for speculative knowledge and its object are identical" (Aristotle, De An. 430a4-6).

The Aristotelian doctrine that understanding and what is understood share their form was accepted by Aquinas. As Brower and Brower-Toland (2008) point out, an adequate interpretation of Aquinas's account—as well as Aristotle's—must tell us what the "intentional possession" of a form by the

mind (or, for us, the possession of a form by a sentence) is. In particular, an interpretation of Aquinas as making a claim of "formal-sameness" must say how forms that are not forms of accidents can occur in the soul in such a way that their occurrence is an accident of the soul. In our context, only forms of something like facts (worldly propositions and sentences) matter; and it seems that facts can occur in the world and also in the mind or in discourse. We think that, in our context, "intentional possession" can be explained in terms of normative incompatibility. Hence, we think that we can avoid the problems that Brower and Brower-Toland identify for the view that they call the "Formal-Sameness Theory." A fuller discussion of such exegetical issues goes beyond the scope of this book.

10 This is the basic idea of information theoretic semantics (see Dretske, 1981). Note, however, that we are developing this idea in a very different direction, which becomes clearer in the next sentences.

11 The traditional notion of substantial form is, of course, intimately connected to many difficult philosophical and exegetical issues. We use the notion in a more metaphysical and abstract sense, and not in the sense of an inner efficient cause of certain properties (see Pasnau, 2004). However, we use the notion merely as an expository device; hence we can ignore the many philosophical and exegetical issues here.

12 Note that one state can be an occurrence of several rational forms, and there might be a way in which an assertion or denial can be an occurrence of several rational forms as well (although we do not discuss this possibility here). According to some theories of substantial forms (notably according to Aquinas's view), this cannot happen with substantial forms. So, this is a potential difference between rational forms and substantial forms.

13 This is again inspired by Aristotle. In the *Categories*, Aristotle notes that assertions and denials are incompatible with each other in a way that mirrors the way in which the states that they represent are incompatible:

> [W]hat underlies an affirmation or negation [is not] itself an affirmation or negation. For an affirmation is an affirmative statement and a negation a negative statement, whereas none of the things underlying an affirmation or negation is a statement. These are, however, said to be opposed to one another as affirmation and negation are; for in these cases, too, the manner of opposition is the same. For in the way an affirmation is opposed to a negation, for example "he is sitting"—"he is not sitting", so are opposed also the actual things underlying each, his sitting—his not sitting. (Aristotle, Cat. 12b5-b16)

We hold that assertion and denial (and acceptance and rejection)—which correspond to Aristotle's affirmation and negation—are opposed to each other normatively. The truth-makers and falsity-makers—which correspond to Aristotle's "things underlying an affirmation or negation"—are opposed to each other alethically. The "manner of opposition" is the same because the two opposition relations are isomorphic.

14 This terminology was introduced by Brandom (2019) in *A Spirit of Trust* (starting at p. 84). However, there Brandom analogizes content and matter,

which we do not do here. As will become clearer in the next chapter, we take conceptual contents to be rational forms.

15 Below, we will discuss whether it can happen that both a truth-maker and a falsity-maker obtain, and whether it can happen that neither obtains. The claim that the first is impossible turns out to be equivalent to the structural rule of Containment, and the claim that the second cannot occur turns out to be equivalent to the structural rule of Cut. Hence, the fact that this passage doesn't rule out either possibility is a reflection of the fact that we are interested in theorizing open reason relations.

16 Quine's inscrutability of reference, as well as Putnam's and Davidson's arguments that truth-conditions do not suffice to determine reference relations can be viewed as versions of this general issue.

17 Fine's notation for fusion is "$\sqcup$", which we will use to express something else in the next chapter. Hence we use "$\Cup$" to denote the fusion of states.

18 Thus the meanings of logical connectives are functions from some worldly propositions to a worldly proposition.

19 We do not foresee any particular problems with the extension of our results to first-order logic. However, there are well-known unresolved questions regarding the truth- and falsity-makers of universal generalizations (Fine 2017c, sec I.7; Jago 2018, chap 5). For example, the truth-maker of "All humans are mammals" cannot just be a fusion of the truth-makers of its actual instances because this fusion does not rule out the possibility that there is another human who is not a mammal. One option is to add totality-facts and to say that a truth-maker of a universal generalization is a fusion of truth-makers for all its instances and the totality-fact that these are all the instances (Armstrong, 2004, 19). Fine prefers a solution in terms of arbitrary objects (Fine, 2017c, 568-569). We suspect that any plausible treatment of these issues can be reproduced within our pragmatics-first approach, so that the isomorphism that is our topic here will continue to hold for the resulting first-order systems.

20 The semantic clause for the falsity-makers of conjunctions that we use here is sometimes called the "inclusive clause." The non-inclusive alternative clause that is sometimes used in truth-maker theory is the following:

(and$-_n$) $s \dashv\!\| B \wedge C$ if and only if $s \dashv\!\| B$ or $s \dashv\!\| C$

This non-inclusive clause allows that a fusion of falsity-makers for the two conjuncts is not an exact falsity-maker of the conjunction. Thus, the state of it neither raining nor being night is not necessarily a falsity-maker of "It is raining and it is night." By contrast, according to the inclusive clause for the falsifiers of conjunctions, a fusion of falsifiers for each conjunct is a falsifier for the conjunction. The dual non-inclusive clause for the truth-makers of disjunctions is the following:

(or$+_n$) $s \Vdash B \vee C$ if and only if $s \Vdash B$ or $s \Vdash C$

As one can see when one considers the correspondence that we will explain below, the non-inclusive clauses correspond to the Ketonen-style sequent rules

of $\mathsf{NMMS}^{\backslash ctr}$. If we adopt the non-inclusive clauses and allow Contraction to fail by defining fusions of states in such a way that fusion is no longer idempotent, that is, the state $s \Cup s$ is not always identical to the state s, then the resulting theory coincides with $\mathsf{NMMS}^{\backslash ctr}$ in its consequence relation, given the same base vocabulary. What this means will become clearer at the end of this chapter, but we will not work out the details. Rather, we will postpone a discussion of non-contractive versions of our theory until the next chapter.

21 We write $\bigwedge\{\gamma_1, ..., \gamma_k\}$ for $\gamma_1 \wedge ... \wedge \gamma_k$, and similarly $\bigvee\{\gamma_1, ..., \gamma_k\}$ for $\gamma_1 \vee ... \vee \gamma_k$.

22 A lot of the work of this assumption can be done by using canonical models (see Fine, 2017a, 647). This works, for instance, for completeness proofs. However, certain points can be seen more clearly with the assumption in place.

23 This formulation differs from Fine's in the quantification over further states u. In the presence of Downward-Closure, the two formulations are equivalent.

24 Stipulating these constraints for atomic sentences suffices (given the semantic clauses) to enforce them for the whole language.

25 Fine takes the relata of consequence to be propositions. We will work with sentences. Assumption 39 ensures that this is unproblematic for our purposes.

26 We will talk about validity in a model, so that the term "validity" has no connotation of a quantification over models in our context. We will later define notions of validity that involve quantification over models. However, this is not essential for the thus defined relations to be relations of validity. We usually use "validity" when we are thinking in model-theoretic terms, even when no quantification over models is involved.

27 The idea behind truth-maker consequence is similar to Bueno and Shalkowski's modalism about logic, which says that B follows from A if and only if the conjunction of A and the negation of B is impossible (Bueno and Shalkowski, 2013, 11-12). In contrast to modalism, however, truth-maker consequence is formulated at a meta-theoretic level.

28 At this point it is again convenient that Definition 38 ensures that we always have a state to evaluate, even if $\Gamma \cup \Delta = \emptyset$ or if Γ has no verifiers or Δ no falsifiers. In such cases the relevant state is $\square$.

29 Indeed, Fine tends to consider it an advantage of his notions of consequence that they don't require us to distinguish the possible from the impossible states.

30 Bob Hale (2020, 132) uses a similar technique in order to capture non-logical necessities in truth-maker theory. And Gil Sagi (2018) has developed a broadly similar idea to fix meanings by constraints on models outside of truth-maker theory. As will become clearer below, we can afford to restrict our models all the way down to a single model. For our models do not settle the truth-values of sentences. Our models are, in this respect, more similar to Kripke frames (without a designated actual world) than to familiar first-order models.

31 These correspondences are explained at a more technical level in other work (Hlobil, 2022a). Here we aim to keep the discussion as accessible as possible. The reader interested in the formal details is referred to this other work and to the appendix to this chapter.

32 Notice that if the norms that govern the use of a sentence are not isomorphic to the modal profile of the worldly proposition that the sentence is meant to represent, then this is not a case of misrepresentation. Rather it is a failure to represent what one means to represent, correctly or incorrectly. Of course, it is an idealization that our conceptual norms are indeed isomorphic to the alethic-modal structure of the worldly propositions that we take them to represent. Thus, we will ultimately have to make room for imperfect or approximate isomorphisms. In this book, however, we ignore such issues, which are closely related to the issue of how our conceptual norms can change rationally over time, and to Semantic Covariant Tracking.

33 See the Appendix for technical details.

34 We also saw that Containment corresponds to Exclusivity. Since rejecting Containment isn't of much importance for us, however, we will only mention it in passing. In the previous chapter, we also talked about Contraction. Giving up Contraction in truth-maker theory would require that we define fusion not as a least upper bound. We will not pursue this here, but we will generalize the insights from the current chapter in the next one, where the issue of Contraction comes up again.

35 This kind of semantics for STT was first developed and presented in (Hlobil, 2022a) and (Hlobil, 2022b). In these papers, there are hints at the greater power of truth-maker semantics relative to the usual strong Kleene semantics for STT. The treatment of nonmonotonic consequence relations in this framework that we are offering below makes good on the promissory notes in these papers.

36 We leave out quantifiers, thus restricting us, in effect, to pure predicate logic (that is, predicate logic without quantifiers or identity). The truth predicate and the name of the liar sentence are the only things we really need from the language of the predicate calculus.

37 This calculus is equivalent, but not identical, to standard formulations of STT in the literature (see Hlobil, 2022a). To see this, note that, as a sequent calculus, STT can be formulated by adding [Lt] and [Rt] to Gentzen's calculus LK. But we know that the narrowly logical part of NMMS is equivalent to Gentzen's LK. So, adding [Lt] and [Rt] to NMMS yields a formulation of STT.

38 Here, we stipulate means of self-reference by fiat: we let λ and $\neg Tr(\bar{\lambda})$ simply be identical. This allows us to avoid the complications of adding self-reference via Gödel numbers. See Ripley (2012, 355) for more details.

39 We will be brief here, as the paradoxes are not really our topic. Readers interested in the technical details are referred to the cited paper.

40 There are well-known relations between ST and the logics LP, K3, and TS (Dicher and Paoli, 2019; Barrio et al., 2015). These relations can be spelled-out in truth-maker theory (see Hlobil, 2022a). If we do that, LP emerges roughly as the logic of impossible falsity-makers. K3 emerges roughly as the logic of possible truth-makers. And TS emerges as the logic that rejects Exclusivity. We will return to these connections in a more general setting in the next chapter.

5 Implication-Space Semantics
Content as Implicational Role

Ulf Hlobil

In the previous chapter, we have seen that the pragmatic-normative understanding of open reason relations from Chapters One to Three is isomorphic to the semantic-representationalist understanding of open reason relations in terms of truth-makers and falsity-makers. In particular, if collections of assertions and denials of atomic sentences are normatively ruled out just in case all fusions of truth-makers of the asserted sentences and falsity-makers of the denied sentences are alethically ruled out, then the pragmatic-normative theory of consequence among sentences from Chapter Three coincides with the semantic-representationalist theory of consequence among worldly propositions from Chapter Four. That is, the consequence relation among sentences defined in pragmatic-normative terms and the consequence relation among worldly propositions, that are represented by these sentences, defined i n semantic-representationalist terms are isomorphic under the mapping of representations to representanda. We thus have a common formal structure of open reason relations that can be explicated in a pragmatic-normative and also in a semantic-representationalist metavocabulary. We called the roles in this formal structure that are shared between sentences and worldly propositions their "rational forms."

In this chapter, we present a theory of these rational forms in themselves and, hence, not as they occur in discursive acts or in worldly states. We thereby abstract away from the different kinds of matter in which rational forms can occur. We present an intrinsic metavocabulary of reason relations, a metavocabulary that allows us to talk about the content of something as its role in reason relations: its role in implications. In other words, we present an account that does not appeal to anything like pragmatic-normative or semantic-representationalist ideas to explain reason relations, but provides rather a metalinguistic codification and explication of reason relations that appeals only to the reason relations

of the base vocabularies themselves. We call this theory "implication-space semantics."[1] From the perspective of implication-space semantics, the pragmatic-normative and the semantic-representationalist theories from previous chapters can be seen as accounts of two particular ways in which rational forms, that is, implicational roles, can be realized.

The chapter is organized as follows. In the first section, we introduce implication-space semantics. We then explain how the pragmatic-normative account in terms of the sequent calculus NMMS and the semantic-representationalist account in terms of truth-makers can be interpreted in implication-space semantics. In the third section, we explain how implication-space semantics reveals the basic structure of reason relations to be that of a monoid defined on a set of pairs, together with a subset of the monoid set. Section 5.4 shows how implication-space semantics can capture non-contractive reason relations and multiplicative additive linear logic. We then present, in Section 5.5, an account of relations among implicational roles, which allows us to recover logics like the logic of paradox (LP) and strong Kleene logic (K3). Section 5.6 concludes.

5.1 Formulating Implication-Space Semantics

The extrinsic metavocabularies of norms governing acceptance and rejection and of states making sentences true or false, which were our topics in the previous two chapters, appealed to something outside of reason relations themselves in order to explain what it means for things to stand in reason relations. We now abstract away from such appeals to external resources and formulate a theory that explicates reason relations merely in terms of the materials already available in reason relations themselves. The theory provides a notion of conceptual contents as roles within reason relations. These conceptual roles are the rational forms shared by contentful sentences and worldly propositions, which we already encountered in the previous chapter.

5.1.1 Implicational Roles

In the previous chapter, we introduced the idea of rational form as the modal role that something plays in relations of exclusion. These relations of exclusion are the generic form of reason relations, of implications (including incompatibilities).[2] Hence, in order to focus on rational forms, we must look directly at roles that things play in reason relations or implications. We will call these roles "implicational roles." We therefore begin by making the notion of an implicational role precise.

Although we want to focus on rational forms and, hence, on implicational roles in abstraction from anything that embodies these forms

or roles, we still need some placeholders for bearers of roles, which we will simply call "bearers." In the two previous chapters, we saw that sentences and worldly propositions can serve as bearers of implicational roles. In this chapter, we abstract away from any particular bearers and focus merely on the roles they play in implications. Our general strategy is to define roles as equivalence classes of bearers or implications among bearers. We will provide a general way to abstract implicational roles from any given set of bearers standing in implication relations, where these implication relations can, crucially, be structurally open reason relations.

We will start by considering an otherwise unspecified set of bearers of implicational roles, and the set of all possible implications among them, which we call "candidate implications." We call the set of all candidate implications the "implication space" of a given set of bearers.

Definition 61 (Role bearers and implication space)**.** The members of a non-empty set $\mathtt{B}$ are bearers of implicational roles. And the set, $\mathbb{S}$, of all the pairs of subsets of $\mathtt{B}$ is the bearer implication space, which contains all the candidate implications among the bearers, so $\mathbb{S} = \mathcal{P}(\mathtt{B}) \times \mathcal{P}(\mathtt{B})$.

We call the first set in a candidate implication its "premises" and the second its "conclusions." Some candidate implications hold and others do not hold, that is, some candidate implications are good implications and some are not. Thus, the actual implications or good implications are a subset of the implication space. We use "implication frame" for a pair of the set of bearers and a subset of good implications.

Definition 62 (Implication and implication frames)**.** Implication is a relation between sets of bearers, $\mathbb{I} \subseteq \mathbb{S}$. An implication frame is a pair $\langle \mathtt{B}, \mathbb{I} \rangle$ of set of bearers, $\mathtt{B}$, and implications among sets of them, $\mathbb{I} \subseteq \mathbb{S}$.

The idea behind this definition is that an open reason relation holds among sets of bearers and that $\Gamma \mid\!\sim \Delta$ is part of this reason relation just in case $\langle \Gamma, \Delta \rangle \in \mathbb{I}$. We will sometimes call the minimal element, $\langle \emptyset, \emptyset \rangle$, of the implication space $\mathtt{e}$, and we denote the maximal element $\langle \mathtt{B}, \mathtt{B} \rangle$ by $\star$. We assume that $\mathtt{e} \notin \mathbb{I}$ but $\star \in \mathbb{I}$.

We want to define the roles that bearers play in implication relations, and these roles should be such that it is possible for distinct bearers to play the same role. To a first approximation, we can think of an implicational role as represented by the collection of bearers that all play the same role. This suggests that we start by asking when two bearers play the same role. Our strategy in this section is to first define an equivalence relation between bearers, and indeed implications and sets thereof, which holds just in case the equivalent items play the same implicational role. And we then define

implicational roles as equivalence classes with respect to this equivalence relation.

In order to define this equivalence relation, consider that, intuitively, the implicational role of a bearer has two parts: firstly, the role that the bearer plays as a premise in implications and, secondly, the role that it plays as a conclusion in implications. We can think of the role that a bearer plays as a premise or a conclusion as the contribution that it makes to the goodness of an implication, as a premise or as a conclusion, respectively. In other words, if we have a bearer, ϕ, of an implicational role, then the role of ϕ as a premise in implications is fully settled by for exactly which sets X and Y we have $\langle X \cup \{\phi\}, Y\rangle \in \mathbb{I}$, that is, by exactly which good implications the bearer figures in as a premise. And the role of ϕ as a conclusion in implications is fully settled by for exactly which sets X and Y we have $\langle X, Y \cup \{\phi\}\rangle \in \mathbb{I}$, that is, by exactly which good implications the bearer figures in as a conclusion. We call the first part of this role of a bearer its "premisory role" and the second part its "conclusory role." The premisory role of a bearer tells us what contribution the bearer makes to good implications as a premise. And the conclusory role of a bearer tells us what contribution the bearer makes to good implications as a conclusion.

Considering premisory and conclusory roles naturally leads to a broadening of our focus from roles of particular bearers to roles of whole implications. To see this, it is helpful to recall the definition of the range of subjunctive robustness from Chapter Three. The idea behind ranges of subjunctive robustness is that, especially in a nonmonotonic setting, it is interesting to consider which additions to an implication do not defeat the implication, if the implication is already good, or turn the implication into a good implication, if it is not already good. If we generalize that definition from $\mid\!\sim$ to $\mathbb{I}$ and also extend the definition to sets of candidate implications, the result is the following:

Definition 63 (Range of subjunctive robustness, $\mathtt{RSR}(\cdot)$). Given $\langle \Gamma, \Delta\rangle \in \mathbb{S}$, its range of subjunctive robustness, $\mathtt{RSR}(\langle \Gamma, \Delta\rangle)$, is the set of pairs, $\langle X, Y\rangle$, such that $\langle \Gamma \cup X, \Delta \cup Y\rangle \in \mathbb{I}$; that is, $\mathtt{RSR}(\langle \Gamma, \Delta\rangle) = \{\langle X, Y\rangle \in \mathbb{S} \mid \langle \Gamma \cup X, \Delta \cup Y\rangle \in \mathbb{I}\}$. The range of subjunctive robustness of a set of candidate implications is the intersection of the ranges of subjunctive robustness of the members; that is, if $H \subseteq \mathbb{S}$, then $\mathtt{RSR}(H) = \{\langle X, Y\rangle \mid \forall \langle \Gamma, \Delta\rangle \in H\,(\langle \Gamma \cup X, \Delta \cup Y\rangle \in \mathbb{I})\}$.

Notice that the set of pairs $\langle X, Y\rangle$ such that $\langle X \cup \{\phi\}, Y\rangle \in \mathbb{I}$ is the range of subjunctive robustness of $\langle \{\phi\}, \emptyset\rangle$. And the set of pairs $\langle X, Y\rangle$ such that $\langle X, Y \cup \{\phi\}\rangle \in \mathbb{I}$ is the range of subjunctive robustness of $\langle \emptyset, \{\phi\}\rangle$.[3] So it follows from what we said above that these ranges of subjunctive robustness fully determine the premisory and conclusory roles of a given

bearer ϕ. This suggests that everything with the same range of subjunctive robustness, as a premise and as a conclusion, plays the same implicational role as the given bearer. Thus, if the two bearers ϕ and ψ are such that the ranges of subjunctive robustness of their occurrences as a premise are identical, that is, if $\mathtt{RSR}\,\langle\{\phi\},\emptyset\rangle = \mathtt{RSR}\,\langle\{\psi\},\emptyset\rangle$,[4] then ϕ and ψ play the same premisory role. And if the ranges of subjunctive robustness of their occurrences as a conclusion are identical, that is, if $\mathtt{RSR}\,\langle\emptyset,\{\phi\}\rangle = \mathtt{RSR}\,\langle\emptyset,\{\psi\}\rangle$, then ϕ and ψ play the same conclusory role.

To formulate this idea in a rigorous way, let us define a notion of equivalence among sets of candidate implications in terms of their ranges of subjunctive robustness. Individual candidate implications, such as $\langle\{\phi\},\emptyset\rangle$ and $\langle\emptyset,\{\phi\}\rangle$, can then be seen as the special case where the sets of candidate implications have only one member.

Definition 64 (Implication equivalence, $\approx$)**.** Let G and F be sets of candidate implications, then $G \approx F$ if and only if $\mathtt{RSR}(G) = \mathtt{RSR}(F)$.

Since it is defined by an identity, this relation is obviously symmetric, reflexive, and transitive. So implication equivalence is indeed an equivalence relation. We can now consider classes of sets of implications that are equivalent to one another, in the sense of implication equivalence. Indeed, we can define implicational roles as such equivalence classes of candidate implications. Given that we treat the roles of bearers in terms of premisory and conclusory roles and we have defined equivalence not only for implications but for sets of implications, we can define implicational roles not only for bearers but also for implications and sets of implications, namely as follows:

Definition 65 (Implicational Role, $\mathcal{R}(\alpha)$)**.** The implicational role, $\mathcal{R}$, of something, α, is its equivalence class under $\approx$, namely[5]:

(1) If $\alpha \subseteq \mathtt{S}$, then $\mathcal{R}(\alpha) = \{x \mid \alpha \approx x\}$.
(2) If $\alpha \in \mathtt{S}$, then $\mathcal{R}(\alpha) = \mathcal{R}(\{\alpha\})$.
(3) If $\alpha \in \mathtt{B}$, then $\mathcal{R}(\alpha) = \langle\mathcal{R}^{+}(\alpha),\mathcal{R}^{-}(\alpha)\rangle$, where $\mathcal{R}^{+}(\alpha) = \mathcal{R}\,\langle\{\alpha\},\emptyset\rangle$ and $\mathcal{R}^{-}(\alpha) = \mathcal{R}\,\langle\emptyset,\{\alpha\}\rangle$.

Clause (3) of this definition says that the implicational role of a bearer, α, is the pair of the implicational role of $\langle\{\alpha\},\emptyset\rangle$ and the implicational role of $\langle\emptyset,\{\alpha\}\rangle$, as defined in clause (2) of the definition. These roles of implications, in turn, are equivalence classes of implications, namely sets of implications that have the same range of subjunctive robustness. Clause (1) gives the general case that defines roles of sets of candidate implications, namely as equivalence classes of such sets.

To illustrate, the premisory role of bearer ϕ is the set of sets of candidate implications with the same range of subjunctive robustness as $\langle\{\phi\},\emptyset\rangle$. For example, the singleton of a candidate implication $\{\langle\Gamma,\Delta\rangle\}$ is in this set if and only if parallel additions of premises or conclusions to $\langle\Gamma,\Delta\rangle$ and to $\langle\{\phi\},\emptyset\rangle$ always either both yield good implications or both yield implications that are not good. In other words, in any candidate implication, we can replace the bearer ϕ as a premise *salva consequentia*—that is, without turning a good implication into a bad one or vice versa—with the combination of Γ as premises and Δ as conclusions, and the other way around. So, if two bearers have the same premisory role, then they can be substituted for each other as premises *salva consequentia*. Conclusory roles are analogous, except that $\langle\{\phi\},\emptyset\rangle$ is changed to $\langle\emptyset,\{\phi\}\rangle$. Hence, if two bearers have the same conclusory role, they can be substituted for each other as conclusions *salva consequentia*. The equivalence classes with respect to ranges of subjunctive robustness that are implicational roles thus capture the idea of "playing the same role in implications."[6]

According to the notion of an implicational role that we have just introduced, what it means to play a particular implicational role is to be a member of a particular equivalence class of things with the same range of subjunctive robustness. For the case of individual bearers of implicational roles, their roles are pairs of such equivalence classes, namely the premisory and the conclusory roles of the bearer.

5.1.2 Conceptual Content and Entailment

An implicational role is an abstraction from the implications among bearers of implicational roles. Two distinct bearers can have the same implicational role. And, there can be roles that are not played by any bearer.[7] For there may be ranges of subjunctive robustness for which there is no particular bearer whose premisory or conclusory role is the equivalence class of that range of subjunctive robustness. With implicational roles, we have entered a realm of abstract entities that we may call "conceptual contents."

What we mean by "conceptual content"—or just "content" for short—is an implicational role that a bearer could play. Thus, a conceptual content is a pair of a premisory role and a conclusory role.

Definition 66 (Conceptual Content)**.** If there are two implicational roles $\mathtt{a}^+ = \mathcal{R}(F)$ and $\mathtt{a}^- = \mathcal{R}(G)$ (for some sets of candidate implications F and G), then the pair of them is a conceptual content, $\mathtt{a}$. That is, $\mathtt{a} = \langle\mathtt{a}^+,\mathtt{a}^-\rangle$, where $\mathtt{a}^+$ is the premisory role and $\mathtt{a}^-$ is the conclusory role of the content. The collection of all contents is called $\mathbb{C}$. (Convention: We will use lowercase typewriter font, as in $\mathtt{a},\mathtt{b},\mathtt{c},\ldots$, for conceptual contents and uppercase typewriter font for sets of contents.)

Notice that two bearers of implicational roles from a given implication space have the same conceptual content if and only if one can always substitute one bearer for the other *salva consequentia*. That is, replacing one of the two bearers by the other as a premise or as a conclusion never turns a good implication into a bad one. We might compare this to possible worlds semantics suggesting that two sentences have the same content just in case they can be substituted for each other *salva veritate* with respect to every possible world, that is, in arbitrary modal contexts.

We can now understand implications among bearers, as the concrete manifestation of a more abstract relation of entailment among conceptual contents, although this concrete manifestation is our only access to the more abstract relation. To see what we mean, it is helpful to define the following operation called "adjunction" on implicational roles.

Definition 67 (Adjunction, $\sqcup$). If $F, G \subseteq \mathbb{S}$, then the adjunction of the implicational roles of F and G, written $\mathcal{R}(F) \sqcup \mathcal{R}(G)$, is $\mathcal{R}(\{\langle \Gamma \cup \Delta, \Theta \cup \Lambda \rangle \mid \langle \Gamma, \Theta \rangle \in F, \langle \Delta, \Lambda \rangle \in G\})$. We write $\bigsqcup_{i=0}^{n} \mathcal{R}_i$ for the adjunction $\mathcal{R}_0 \sqcup \ldots \sqcup \mathcal{R}_n$.

Adjunction combines implicational roles by combining by set-theoretic union the premises and the conclusions of the candidate implications that play these roles. This yields a unique result for different members of the implicational roles, so that if $\mathcal{R}(F) = \mathcal{R}(F')$, then $\mathcal{R}(F) \sqcup \mathcal{R}(G) = \mathcal{R}(F') \sqcup \mathcal{R}(G)$ (see Appendix, Proposition 105).

We can use this notion of adjunction to define an entailment relation among contents in terms of implications among bearers. The idea is that the entailment relation holds among sets of contents $\mathtt{G}$ and $\mathtt{D}$ just in case, if we take as premises bearers that play the premisory implicational roles of the contents in $\mathtt{G}$ and we take as conclusions bearers that play the conclusory implicational roles of the contents in $\mathtt{D}$, then this always yields a good implication. More precisely, we say that a set of contents, $\mathtt{G}$, entails a set of contents, $\mathtt{D}$, just in case every set of candidate implications that plays the implicational role of the adjunction of the premisory roles of the elements of $\mathtt{G}$ and the conclusory roles of the elements of $\mathtt{D}$ contains only good implications.

Definition 68 (Content entailment in an implication frame, $\Vdash\!\sim$). Given an implication frame $\langle \mathtt{B}, \mathbb{I} \rangle$, let $\mathtt{G}, \mathtt{D} \subseteq \mathbb{C}$ and $\mathtt{G} = \{\mathtt{g}_0, \ldots, \mathtt{g}_n\}$ and $\mathtt{D} = \{\mathtt{d}_0, \ldots, \mathtt{d}_m\}$, then $\mathtt{G} \Vdash\!\sim \mathtt{D}$ holds in the implication frame if and only if $\bigcup(\bigsqcup_{i=0}^{n} \mathtt{g}_i^+ \sqcup \bigsqcup_{j=0}^{m} \mathtt{d}_j^-) \subseteq \mathbb{I}$.

To see what this definition means, let us consider an example. Suppose we have two sets of bearers of implicational roles $\Gamma = \{g_0, ..., g_n\}$ and $\Delta = \{d_0, ..., d_m\}$. Their implicational roles are $\mathcal{R}(g_i) = \langle \mathtt{g}_i^+, \mathtt{g}_i^- \rangle$ and $\mathcal{R}(d_j) = \langle \mathtt{d}_j^+, \mathtt{d}_j^- \rangle$, for all $0 \leq i \leq n$ and $0 \leq j \leq m$. The contents of the bearers in Γ entail the contents of the bearers in Δ if and only if the adjunction of the premisory roles of the bearers in Γ and the conclusory roles of the bearers in Δ has only members that are good implications. In other words, we have to check the adjunction of the following roles: $\mathcal{R}\langle\{g_0\}, \emptyset\rangle = \mathtt{g}_0^+$,..., $\mathcal{R}\langle\{g_n\}, \emptyset\rangle = \mathtt{g}_n^+$, and $\mathcal{R}\langle\emptyset, \{d_0\}\rangle = \mathtt{d}_0^-$, ..., and $\mathcal{R}\langle\emptyset, \{d_m\}\rangle = \mathtt{d}_m^-$. When we perform this adjunction, the result is: $\mathcal{R}\langle\Gamma, \Delta\rangle$. Hence, the entailment holds if and only if every set of candidate implications whose range of subjunctive robustness is $\mathtt{RSR}\langle\Gamma, \Delta\rangle$ includes only good implications.

We have now seen how we can define conceptual content and entailment for any given implication frame, $\langle \mathtt{B}, \mathbb{I} \rangle$.[8] Contents and the entailment relation among them are defined by abstracting away from implications among concrete bearers, in terms of abstract implicational roles.

5.1.3 Models and Interpretations

The notions of content and entailment that we have just introduced can be used to interpret sentences of a language. Recall that we have not specified what bearers of implicational roles are. So, if we are given an implication frame that uses non-linguistic bearers of implicational roles or bearers from one language, we can use the contents defined by that implication frame as interpretants of a(nother) language. As we shall see in this subsection, we can then use such interpretations to give a treatment of logical vocabulary.

In order to give interpretations of logical vocabulary, however, we first introduce another operation on implicational roles, which we call "symjunction." It is another way to combine implicational roles, in addition to adjunction.

Definition 69 (Symjunction, $\sqcap$)**.** Let $F, G \subseteq \mathbb{S}$, then: $\mathcal{R}(F) \sqcap \mathcal{R}(G) = \mathcal{R}(F \cup G)$.

Similarly to adjunction, the symjunction of roles is unique, in the sense that it does not depend on which bearers we use to pick out the roles whose symjunction we compute. That is, if $\mathcal{R}(F) = \mathcal{R}(F')$, then $\mathcal{R}(F) \sqcap \mathcal{R}(G) = \mathcal{R}(F') \sqcap \mathcal{R}(G)$ (see Appendix, Proposition 106).

Although symjunction is defined in terms of a union of sets of candidate implications, we can think of it as, in a way, isolating the part that two implicational roles share. To see this, note that $\mathcal{R}(F \cup G)$ is the set of sets of candidate implications whose range of subjunctive robustness is $\mathtt{RSR}(F \cup$

G). Since we defined the range of subjunctive robustness of a set, however, as the intersection of the ranges of subjunctive robustness of its members, $\mathtt{RSR}(F \cup G)$ is the same as $\bigcap_{x \in F \cup G} \mathtt{RSR}(x)$, which is $\mathtt{RSR}(F) \cap \mathtt{RSR}(G)$. The implicational role of a union of sets of candidate implications is, therefore, the sets of candidate implications that share their range of subjunctive robustness with the intersection of the ranges of subjunctive robustness of the sets of which we take the union. Hence, all of the things that yield a good implication when combined with something that plays the role $\mathcal{R}(F)$, and also when combined with something that plays the role $\mathcal{R}(G)$, yield a good implication when combined with something that plays the role of their symjunction $\mathcal{R}(F) \sqcap \mathcal{R}(G)$. In this sense, as the symbols we have chosen for them indicate, adjunction and symjunction operate on conceptual roles in ways somewhat analogous to the way union and intersection operate on sets.

We can now define an interpretation of a language as a function that assigns each sentence a conceptual content.

Definition 70 (Interpretation Function, $[\![\cdot]\!]$). An interpretation function $[\![\cdot]\!]$ maps sentences of a language $\mathfrak{L}$ to some conceptual contents $\mathbb{C}$. If $A \in \mathfrak{L}$ is an atomic sentence, then $[\![A]\!] =_{df.} \langle \mathtt{a}^+, \mathtt{a}^- \rangle \in \mathbb{C}$. The logical-connective clauses that an interpretation must respect are as follows:

$[\![\neg A]\!] =_{df.} \langle \mathtt{a}^-, \mathtt{a}^+ \rangle$,
$[\![A \to B]\!] =_{df.} \langle \mathtt{a}^- \sqcap \mathtt{b}^+ \sqcap (\mathtt{a}^- \sqcup \mathtt{b}^+), \mathtt{a}^+ \sqcup \mathtt{b}^- \rangle$,
$[\![A \wedge B]\!] =_{df.} \langle \mathtt{a}^+ \sqcup \mathtt{b}^+, \mathtt{a}^- \sqcap \mathtt{b}^- \sqcap (\mathtt{a}^- \sqcup \mathtt{b}^-) \rangle$,
$[\![A \vee B]\!] =_{df.} \langle \mathtt{a}^+ \sqcap \mathtt{b}^+ \sqcap (\mathtt{a}^+ \sqcup \mathtt{b}^+), \mathtt{a}^- \sqcup \mathtt{b}^- \rangle$.

Interpretations of sets of sentences are the set of the interpretants of the sentences, that is, $[\![\Gamma]\!] = \{[\![\gamma]\!] \mid \gamma \in \Gamma\}$.

In general, these semantic clauses assign contents to logically complex sentences by assigning ranges of subjunctive robustness to their use as premises and as conclusions. The clauses place logically complex sentences in an equivalence class of sets of implications with the same range of subjunctive robustness. That makes sense because, according to implication-space semantics, what it is for a sentence to have a content is for it to be a member of an equivalence class with respect to ranges of subjunctive robustness.

To understand how interpretation functions work, it may be helpful to notice that negation swaps the premisory and conclusory roles, so that the content that an interpretation function assigns to a negation $\neg\phi$ is always the content that is like the content that the function assigns to ϕ except that the premisory and the conclusory roles are swapped. Hence, the ranges of subjunctive robustness of $\langle \{\neg\phi\}, \emptyset \rangle$ and $\langle \emptyset, \{\neg\phi\} \rangle$ are, respectively, the

ranges of subjunctive robustness of $\langle \emptyset, \{\phi\} \rangle$ and $\langle \{\phi\}, \emptyset \rangle$. So, whatever yields a good implication when combined with $\langle \{\neg\phi\}, \emptyset \rangle$ also yields a good implication when combined with $\langle \emptyset, \{\phi\} \rangle$. This is how negation allows us to make explicit incompatibility. For we encode the fact that a set of bearers Γ is incompatible with the bearer ϕ in our implications as $\langle \Gamma \cup \{\phi\}, \emptyset \rangle \in \mathbb{I}$. And our negation clauses ensure that we can substitute the negation for the negatum on the other side *salva consequentia* to get: $\langle \Gamma, \{\neg\phi\} \rangle \in \mathbb{I}$. Thus, the negation makes explicit, on the right side of an implication, that it is ruled out by what is on the left side. This is precisely the explicative potential that we described as the essential function of negation in Chapter Three.

For the interpretation of conditionals, the content of a conditional always has as its conclusory role the adjunction of the premisory role of the antecedent and the conclusory role of the consequent. That means that the range of subjunctive robustness of $\langle \Gamma, \Delta \cup \{\phi \to \psi\} \rangle$ is always identical to the range of subjunctive robustness of $\langle \Gamma \cup \{\phi\}, \Delta \cup \{\psi\} \rangle$. This ensures that the conditional makes implications explicit in the way described in Chapter Three. Namely, it ensures that $\langle \Gamma \cup \{\phi\}, \Delta \cup \{\psi\} \rangle \in \mathbb{I}$ just in case $\langle \Gamma, \Delta \cup \{\phi \to \psi\} \rangle \in \mathbb{I}$, which is a formulation of the Deduction-Detachment Condition on Conditionals from Chapter Three.

And the premisory role that an interpretation function assigns to a conditional is the symjunction of the conclusory role of the antecedent, the premisory role of the consequent, and the adjunction of these roles. That means that the range of subjunctive robustness of $\langle \Gamma \cup \{\phi \to \psi\}, \Delta \rangle$ is always the intersection of the range of subjunctive robustness of $\langle \Gamma, \Delta \cup \{\phi\} \rangle$ and the range of subjunctive robustness of $\langle \Gamma \cup \{\psi\}, \Delta \rangle$ and the range of subjunctive robustness of $\langle \Gamma \cup \{\psi\}, \Delta \cup \{\phi\} \rangle$. In this way, the semantic clauses for the conditional assign a conceptual content to conditionals by placing the use of the conditional as a premise and as a conclusion in equivalence classes with respect to ranges of subjunctive robustness, that is, by assigning the conditional its premisory and conclusory implicational role. The clauses for conjunction and disjunction can be understood in an analogous way.

We can now define models of implication-space semantics. In general, a model is a space of semantic interpretants together with an interpretation that assigns an interpretant from that space to every sentence of a given language. This suggests the following definition of models.

Definition 71 (Models)**.** A model, $\mathcal{M}$, is a pair $\left\langle \mathbb{C}, [\![\cdot]\!]^{\mathcal{M}} \right\rangle$ consisting of a set of contents, $\mathbb{C}$, and an interpretation function $[\![\cdot]\!]^{\mathcal{M}}$ that maps all sentences of a given language to contents in $\mathbb{C}$.

We haven't used the idea of truth anywhere in our construction in this chapter, and we also have no use for the traditional notion of truth in a model. Our basic semantic notion is not truth but implication, that is, reason relations. Accordingly, we can define a notion of implication among sets of sentences in a model.

Definition 72 (Implication in a model, $\overset{\mathcal{M}}{\mid\!\sim}$)**.** We say that the sentences Γ model-theoretically imply the sentences Δ in model $\mathcal{M}$, written $\Gamma \overset{\mathcal{M}}{\mid\!\sim} \Delta$, if and only if the corresponding entailment holds among their interpretants, $\mathtt{G} = [\![\Gamma]\!]^{\mathcal{M}}$ and $\mathtt{D} = [\![\Delta]\!]^{\mathcal{M}}$, in that model, that is, if and only if $\mathtt{G} \Vdash\!\sim \mathtt{D}$.

We generalize this idea from a single model to arbitrary sets of models as follows:

Definition 73 (Implication in m-models, $\overset{m}{\mid\!\sim}$)**.** Given a subset, m, of all implication-space models, for a given language, we say that the sentences in Γ model-theoretically imply the sentences in Δ in m, written $\Gamma \overset{m}{\mid\!\sim} \Delta$, if and only if $\Gamma \overset{\mathcal{M}}{\mid\!\sim} \Delta$ for all $\mathcal{M} \in m$.

We are now ready to explore the effects of interpreting sentences by assigning contents to them. A first and easy thing to note is that if we abstract contents from an implication frame whose implication space is given by a language, then interpreting the sentences of that language in terms of the abstracted contents in the most straightforward way yields the implication relation of the implication frame itself.

Proposition 74. *Let* $\mathbb{C}$ *be the conceptual contents abstracted from the implication frame,* $\langle \mathfrak{L}, \mathbb{I} \rangle$ *and, hence,* $\mathbb{S} = \mathcal{P}(\mathfrak{L}) \times \mathcal{P}(\mathfrak{L})$*, and let* $\mathcal{M}$ *be the pair of* $\mathbb{C}$ *and an interpretation function such that* $[\![\phi]\!]^{\mathcal{M}} = \mathcal{R}(\phi)$ *for all* $\phi \in \mathfrak{L}$*. Then* $\langle \Gamma, \Delta \rangle \in \mathbb{I}$ *if and only if* $\Gamma \overset{\mathcal{M}}{\mid\!\sim} \Delta$*.*

Proof. Suppose that $\langle \Gamma, \Delta \rangle \in \mathbb{I}$. Then, $\langle \Gamma, \Delta \rangle \approx F$ only if $\langle \emptyset, \emptyset \rangle \in \mathtt{RSR}(F)$ and, hence, $F \subseteq \mathbb{I}$. Hence, $\bigcup(\mathcal{R}\langle \Gamma, \Delta \rangle) \subseteq \mathbb{I}$. So, $[\![\Gamma]\!]^{\mathcal{M}} \Vdash\!\sim [\![\Delta]\!]^{\mathcal{M}}$. For the other direction, suppose that $[\![\Gamma]\!]^{\mathcal{M}} \Vdash\!\sim [\![\Delta]\!]^{\mathcal{M}}$. Then, $\bigcup(\mathcal{R}\langle \Gamma, \Delta \rangle) \subseteq \mathbb{I}$. But $\langle \Gamma, \Delta \rangle \in \mathcal{R}\langle \Gamma, \Delta \rangle$. Hence, $\langle \Gamma, \Delta \rangle \in \mathbb{I}$. ■

This proposition tells us that we can always abstract conceptual contents from an implication frame and then use them to interpret the bearers of that very implication frame, in an, as it were, homophonic way, while leaving the reason relation among the bearers of implicational roles unchanged. Thus, every implication frame interprets itself. Since every vocabulary is an implication frame (namely a lexicon and a consequence relation over it),

it follows that every vocabulary interprets itself. That is reassuring, but it is of course not very illuminating. As we will see in the next subsection, however, conceptual contents can also be used to interpret the theories from the previous chapters in a way that lets us focus on their isomorphic implicational roles, that is, on rational forms.

5.2 Interpreting NMMS and Truth-Maker Theory

In this section, we show that the pragmatic-normative theory formulated in the language of the sequent calculus from Chapter Three and the semantic-representationalist truth-maker theory from Chapter Four can be interpreted in implication-space semantics. Both theories can be understood as having the structure of an implication-space semantics, each with its own particular extrinsic understanding of the bearers and exclusion relations that explain conceptual contents. We start with the sequent calculus for the pragmatic-normative theory and then turn to truth-maker theory.

5.2.1 Interpreting Normative Bilateralism

As we have seen in the previous subsection, we can start with an implication frame, $\langle \mathbb{B}, \mathbb{I} \rangle$, abstract from it its contents, $\mathbb{C}$, and then use those contents to interpret a given language, in the sense of a set of uninterpreted sentences. Now, if the language under consideration is not simply a set of uninterpreted sentences but rather a vocabulary in the sense we introduced in earlier chapters, then the reason relations over the sentences put restrictions on viable interpretations. In particular, if we start with a base vocabulary, like those of Chapter Three, then the base consequence relation puts constraints on viable interpretations of the language. Let us make this thought precise.

Recall from Chapter Three that a base, $\mathfrak{B}$, is a vocabulary that consists of just atomic sentences $\mathfrak{L}_{\mathfrak{B}}$ and an implication relation, $\mid\!\sim_{\mathfrak{B}}$, among sets of these atomic sentences. In order to find a class of models that can capture the implication relation of such a base, we must look at models that interpret the sentences of $\mathfrak{L}_{\mathfrak{B}}$ in an appropriate way. We say that such models are fit for the base at issue.

Definition 75 (Model fitness for base)**.** A model, $\mathcal{M} = \left\langle \mathbb{C}, [\![\cdot]\!]^{\mathcal{M}} \right\rangle$ is fit for a base $\mathfrak{B} = \left\langle \mathfrak{L}_{\mathfrak{B}}, \mid\!\sim_{\mathfrak{B}} \right\rangle$ if and only if, for all $\langle \Gamma, \Delta \rangle \in \mid\!\sim_{\mathfrak{B}}$, the model is such that $[\![\Gamma]\!]^{\mathcal{M}} \Vdash\!\sim [\![\Delta]\!]^{\mathcal{M}}$.

If b is the set of models that are fit for the base $\mathfrak{B}$, then the implications in b-models ($\overset{b}{\mid\!\sim}$) are the implications that hold in all models that are fit for base $\mathfrak{B}$. If we restrict ourselves to models that are fit for the base $\mathfrak{B}$, then the consequence relation defined by the sequent calculus $\mathsf{NMMS}_{\mathfrak{B}}$ and the model-theoretic implication relation coincide.

Theorem 76. *For any base vocabulary* $\mathfrak{B} = \langle \mathfrak{L}_{\mathfrak{B}}, \underset{\mathfrak{B}}{\mid\!\sim} \rangle$ *and sentences* $\Gamma, \Delta \subseteq \mathfrak{L}$ *in the logically extended lexicon of the base,* $\Gamma \overset{b}{\mid\!\sim} \Delta$ *if and only if* $\Gamma \succ \Delta$ *is derivable in* $\mathsf{NMMS}_{\mathfrak{B}}$. *(Appendix, Theorem 109)*

This theorem immediately entails that, relative to any base vocabulary, the sequent calculus NMMS is sound and complete with respect to the implication-space semantics developed in this chapter. To see why this theorem holds (without going through the details of the proof in the Appendix), it is helpful to notice that, firstly, the match between atomic sequents and the implication-space models is ensured by the restrictions to models that are fit for the base at issue. Secondly, the semantic clauses for the logical constants in implication-space semantics correspond to the sequent rules of NMMS in a perhaps surprisingly direct way. To see this consider the rules for the conditional, as an example:

$$\frac{\Gamma \succ \Delta, A \qquad B, \Gamma \succ \Delta \qquad B, \Gamma \succ \Delta, A}{\Gamma, A \to B \succ \Delta}\,[\mathrm{L}\to] \qquad \frac{\Gamma, A \succ B, \Delta}{\Gamma \succ A \to B, \Delta}\,[\mathrm{R}\to]$$

Since these rules are invertible, as we have shown in Chapter Three, the right rule implies that $\Gamma \succ A \to B, \Delta$ is derivable if and only if $\Gamma, A \succ B, \Delta$ is derivable. This is the Deduction-Detachment Condition from Chapter Three. It corresponds in implication-space semantics to the claim that $\Gamma \overset{\mathcal{M}}{\mid\!\sim} A \to B, \Delta$ holds in a model if and only if $\Gamma, A \overset{\mathcal{M}}{\mid\!\sim} B, \Delta$ holds in that model. Now, $\Gamma \overset{\mathcal{M}}{\mid\!\sim} A \to B, \Delta$ holds just in case any implication that plays the role of $\langle \Gamma, \Delta \rangle$ yields a good implication when it is combined with anything that plays the conclusory role of $A \to B$. So $\{\langle \Gamma, \Delta \rangle\}$ must be in the range of subjunctive robustness of $\langle \emptyset, \{A \to B\} \rangle$, and the same must hold for any other implications with the same role as $\langle \Gamma, \Delta \rangle$. Our semantic clauses for the conditional tell us that the role of $\langle \emptyset, \{A \to B\} \rangle$ is $\mathtt{a}^+ \sqcup \mathtt{b}^-$, that is, it is the role of $\langle \{A\}, \{B\} \rangle$. But two things with the same role share their ranges of subjunctive robustness. So everything with the role $\{\langle \Gamma, \Delta \rangle\}$ is in the range of subjunctive robustness of $\langle \emptyset, \{A \to B\} \rangle$ just in case it is in the range of subjunctive robustness of $\langle \{A\}, \{B\} \rangle$. So, $\Gamma, A \overset{\mathcal{M}}{\mid\!\sim} B, \Delta$.

For the left rule, [L→], the analogous relation holds. Because the rule is invertible, it tells us that $\Gamma, A \to B \succ \Delta$ is derivable if and only if all of

$\Gamma \succ A, \Delta$ and $\Gamma, B \succ \Delta$ and $\Gamma, B \succ A, \Delta$ are derivable. Suppose that $\Gamma, A \to B \overset{\mathcal{M}}{\mid\!\sim} \Delta$. Then the premisory role of $A \to B$ is such that combining anything that plays that role and anything with the role of $\langle \Gamma, \Delta \rangle$ yields a good implication. Hence, $\langle \Gamma, \Delta \rangle$ (and anything that plays its role) is in the range of subjunctive robustness of $\langle \{A \to B\}, \emptyset \rangle$. The semantic clause for the conditional informs us that the role of $\langle \{A \to B\}, \emptyset \rangle$ is $\mathtt{a}^- \sqcap \mathtt{b}^+ \sqcap (\mathtt{a}^- \sqcup \mathtt{b}^+)$. And, as we noted above, the range of subjunctive robustness of a symjunction of roles is the intersection of the ranges of subjunctive robustness of the roles whose symjunction it is. So, the range of subjunctive robustness of $\mathtt{a}^- \sqcap \mathtt{b}^+ \sqcap (\mathtt{a}^- \sqcup \mathtt{b}^+)$ is the intersection of the ranges of subjunctive robustness of $\langle \emptyset, \{A\} \rangle$ and $\langle \{B\}, \emptyset \rangle$ and $\langle \{B\}, \{A\} \rangle$. Therefore, $\langle \Gamma, \Delta \rangle$ (and anything that plays its role) is in the ranges of subjunctive robustness of all of $\langle \emptyset, \{A\} \rangle$ and $\langle \{B\}, \emptyset \rangle$ and $\langle \{B\}, \{A\} \rangle$ (and anything that plays the same roles). And that is just another way to say that $\Gamma \overset{\mathcal{M}}{\mid\!\sim} A, \Delta$ and $\Gamma, B \overset{\mathcal{M}}{\mid\!\sim} \Delta$ and $\Gamma, B \overset{\mathcal{M}}{\mid\!\sim} A, \Delta$ all hold.

The other sequent rules of NMMS correspond to the other semantic clauses of implication-space semantics in an analogous way (see Appendix, Proposition 108). So, the rules of NMMS are not only equivalent to the semantic clauses of truth-maker theory, as we have shown in the previous chapter, but they are also equivalent to the semantic clauses of implication-space semantics. Indeed, we can formulate this correspondence in a general way as follows. The first element in the roles defined by the semantic clauses corresponds to the left rule in the sequent calculus, and the second element corresponds to the right rule in the sequent calculus. The roles super-scripted with a "+" stem from sentences that occur on the left in a top sequent, and the roles super-scripted with a "−" stem from sentences that occur on the right in a top sequent. An adjunction indicates that the adjoined roles stem from sentences in a single top sequent. And a symjunction indicates that the symjoined roles stem from sentences that occur in different top sequents. Given that the contexts are always shared in all the sequents of any rule application, using this correspondence, the semantic clauses above uniquely determine the sequent rules of NMMS, and the other way around.

All the virtues and features of NMMS carry over to $\overset{b}{\mid\!\sim}$. In particular, the logical vocabulary of implication-space semantics makes explicit reason relations in the same sense as NMMS. The Deduction-Detachment Condition on Conditionals, the Incoherence-Incompatibility Condition on Negation, the Antecedent-Adjunction Condition on Conjunctions, and the Succedent-Summation Condition on Disjunctions all hold in every implication-space model. Moreover, $\overset{b}{\mid\!\sim}$ can codify nonmonotonic and nontransitive implication relations; thus, it allows us to codify and theorize

open reason relations. Similarly, the relation between NMMS and classical logic carries over to implication-space semantics. Let us spell out the last point in a bit more detail.

We know from our discussion of NMMS that if we start with a base consequence relation that includes all and only the instances of Containment, then the logically extended consequence relation coincides with classical propositional consequence. And if the base consequence relation includes all instances of Containment, then the logically extended consequence relation includes the classical propositional consequence relation. Hence, the following is an immediate corollary of Theorem 76.

Corollary 77. *If the base* $\mathfrak{B} = \left\langle \mathfrak{L}_{\mathfrak{B}}, \mathrel{|\!\sim}_{\mathfrak{B}} \right\rangle$ *is such that, for all* $\Gamma, \Delta \subseteq \mathfrak{L}_{\mathfrak{B}}$, *we have* $\Gamma \mathrel{|\!\sim}_{\mathfrak{B}} \Delta$ *just in case* $\Gamma \cap \Delta \neq \emptyset$, *then* $\models_{CL} = \overset{b}{\mathrel{|\!\sim}}$. *And if the base is such that, for all* $\Gamma, \Delta \subseteq \mathfrak{L}_{\mathfrak{B}}$, *we have* $\Gamma \mathrel{|\!\sim}_{\mathfrak{B}} \Delta$ *if* $\Gamma \cap \Delta \neq \emptyset$, *then* $\models_{CL} \subseteq \overset{b}{\mathrel{|\!\sim}}$.

Just as in NMMS, we can say that, for bases that obey Containment, there is a narrowly logical part of the consequence relation defined by implication-space semantics, and this narrowly logical part is simply classical logic. Given any base that obeys Containment, its narrowly logical part is the part of the logically extended consequence relation that it shares with the extensions of all other bases that obey Containment. We could also formulate this connection thus: If we say that a proper conceptual content is a conceptual content, $\mathtt{p}$, such that, for all $\mathtt{q}$, $\bigcup(\mathtt{p}^+ \sqcup \mathtt{p}^- \sqcup \mathtt{q}) \subseteq \mathbb{I}$, then the implications in any implication-space model with only proper contents include all implications of classical logic. It is supraclassical.

5.2.2 Interpreting Truth-Maker Theory

We now turn to truth-maker theory. We can define an implication frame for any modalized state space, $\langle S, S^{\diamond}, \sqsubseteq \rangle$, by letting the bearers be worldly propositions, that is, pairs of sets of states from S, and defining the good implications by appeal to impossible states in the way we did in the previous chapter.

Definition 78 (Implication frame of a modalized state space)**.** Given a modalized state space, $\langle S, S^{\diamond}, \sqsubseteq \rangle$, its implication frame is $\langle \mathtt{B}, \mathbb{I}_{\overline{\diamond}} \rangle$, with $\mathtt{B} = \mathcal{P}(S) \times \mathcal{P}(S)$ being the set of worldly propositions of the state space, which we denote by $p = \langle p^+, p^- \rangle$ and the like. And $\mathbb{I}_{\overline{\diamond}} \subseteq \mathcal{P}(\mathtt{B}) \times \mathcal{P}(\mathtt{B})$ is the relation such that $\langle \{p_1, ..., p_n\}, \{q_1, ..., q_m\} \rangle \in \mathbb{I}_{\overline{\diamond}}$ if and only if $s \notin S^{\diamond}$, for all states $s = t_1 \Cup ... \Cup t_n \Cup u_1 \Cup ... \Cup u_m$ such that $t_i \in p_i^+$, for all $1 \leq i \leq n$, and $u_j \in q_j^-$, for all $1 \leq j \leq m$.

This definition is long but all it does is to take a modalized state space and return the set of worldly propositions in that state space and the implication relation among these propositions, according to the idea of truth-maker consequence that we introduced in the previous chapter. According to this idea, a set of worldly propositions implies another just in case any fusion of a truth-maker of each proposition in the first set and a falsity-maker of each proposition in the second set is an impossible state. With p^+ being the set of truth-makers and p^- being the set of falsity-makers of a worldly proposition p, this conception of truth-maker consequence yields the implication relation $\mathbb{I}_{\overline{\diamond}}$. So, the implication frame of a modalized state space is really just the set of worldly propositions in the state space of truth-maker theory and the consequence relation among them, according to truth-maker consequence.

We can now interpret a modalized state space by its implication frame. We think of the modalized state space as providing us with a "vocabulary," that is, a "lexicon" and a consequence relation over that "lexicon" (although worldly propositions are, of course, not sentences). We then let this vocabulary interpret itself, as we did in Proposition 74 above. This yields an interpretation of truth-maker theory in implication-space semantics.

Theorem 79. *Let $\mathcal{M}$ be the implication-space model defined by the implication frame, $\langle \mathtt{B}, \mathbb{I}_{\overline{\diamond}} \rangle$, of the modalized state space $\langle S, S^{\diamond}, \sqsubseteq \rangle$ and the interpretation function such that $[\![x]\!]^{\mathcal{M}} = \mathcal{R}(x)$ for all $x \in \mathtt{B}$. Then $P \Vdash_{PI} C$ holds in the modalized state space if and only if $P \overset{\mathcal{M}}{\mid\!\sim} C$. (Appendix, Theorem 110)*

This theorem says that, for any modalized state space, we can interpret its worldly propositions by assigning them contents from implication frames in a way that preserves the implication relations among these worldly propositions. Notice that the theorem does not make any use of a language in the usual sense. What we interpret is a set of worldly propositions.[9]

Notice that there seems to be a certain asymmetry between implication-space semantics and truth-maker theory in Theorem 79. For one might wonder why we cannot map the implication relation among contents in truth-maker theory, $\Vdash_{PI}$, directly to the entailment relation among contents in implication-space semantics, $\Vdash\!\sim$, without invoking the middle man of an interpretation in $\overset{\mathcal{M}}{\mid\!\sim}$. The answer is that truth-maker theory allows for two distinct worldly propositions, A and B, such that, for all sets of worldly propositions P and C, we have $P, A \Vdash_{PI} C$ just in case $P, B \Vdash_{PI} C$ and $P \Vdash_{PI} A, C$ just in case $P \Vdash_{PI} B, C$. That is, there can be two distinct worldly propositions, in a truth-maker model, that play the same roles

in propositional implications. In particular, these are worldly propositions whose truth-makers and falsity-makers make the same contributions, respectively, to the impossibility or possibility of states of which they are parts. The analogous situation, however, cannot occur in implication-space semantics. That is, if, for all sets of conceptual contents $\mathtt{G}$ and $\mathtt{D}$, we have $\mathtt{G}, \mathtt{p} \Vdash\!\sim \mathtt{D}$ just in case $\mathtt{G}, \mathtt{q} \Vdash\!\sim \mathtt{D}$ and $\mathtt{G} \Vdash\!\sim \mathtt{p}, \mathtt{D}$ just in case $\mathtt{G} \Vdash\!\sim \mathtt{q}, \mathtt{D}$, then $\mathtt{p} = \mathtt{q}$. After all, a conceptual content is an equivalence class with respect to roles in good implications. So contents that play the same such roles are identical. It is only at the level of content bearers that we can have distinct items that play the same role in implications. Hence, to map a relation in implication-space semantics one-to-one to a consequence relation in a truth-maker model, we must use a space of content bearers and not conceptual contents themselves. In this sense, the worldly propositions of truth-maker theory are more like sentences than the conceptual contents of implication-space semantics: there can be distinct worldly propositions with the same content, but there cannot be distinct contents with the same content.

If we use interpretation functions on both sides, in truth-maker theory and in implication-space semantics, we can avoid this asymmetry. Indeed, it follows from Theorem 79 that if there is a truth-maker model in which exactly a particular set of implications hold among the interpreted sentences, then there is an implication-space model in which exactly the same implications hold. The theorem ensures that for every truth-maker model, there is a parallel implication frame model such that the consequence relation defined by these models coincide. To make this precise, let us define explicitly what we mean by "parallel models."

Definition 80 (Parallel Models)**.** Let $\mathfrak{L}$ be a language, then an implication-space model $\mathcal{M}$ and a truth-maker model $\mathcal{M}'$ are parallel just in case, for all $\phi \in \mathfrak{L}$, $[\![\phi]\!]^{\mathcal{M}} = |\phi|^{\mathcal{M}'}$ and the implication frame of $\mathcal{M}$ is the implication frame of the modalized state space of $\mathcal{M}'$.[10]

Truth-maker models and implication-space models that are parallel in this sense codify the same consequence relation over the common language that they interpret.

Proposition 81. *If the implication-space model $\mathcal{M}$ and truth-maker model $\mathcal{M}'$ are parallel, then $\Gamma \overset{\mathcal{M}}{\Vdash\!\sim} \Delta$ just in case $\Gamma \underset{TM}{\models\!\!=} \Delta$ in $\mathcal{M}'$. (Appendix, Theorem 111)*

If we are given any truth-maker model, we can construct the parallel implication-space model. We simply take the implication frame of the modalized state space of the truth-maker model and add the interpretation

such that $[\![\phi]\!]^{\mathcal{M}} = |\phi|^{\mathcal{M}'}$ for all $\phi \in \mathfrak{L}$. Hence, it follows from the last theorem that if there is a truth-maker model that codifies a particular consequence relation over a language, then there is an implication-space model that codifies the same consequence relation over that language.

We can also go in the other direction. If we are given an implication-space model that codifies a particular consequence relation over a language, then we can construct a truth-maker model that codifies the same consequence relation over that language.

Proposition 82. *If there is an implication-space model $\mathcal{M}$ such that $\overset{\mathcal{M}}{\mid\!\sim} = X$, then there is a truth-maker model $\mathcal{M}'$ such that $\overline{\models_{TM}} = X$. (Appendix, Proposition 112)*

Putting the two directions together yields:

Theorem 83. *There is an implication-space model $\mathcal{M}$ such that $\overset{\mathcal{M}}{\mid\!\sim} = X$ if and only if there is a truth-maker model $\mathcal{M}'$ such that $\overline{\models_{TM}} = X$.*

It follows immediately from this theorem that truth-maker theory and implication-space semantics are equivalent in their power to provide counterexamples to implications. Hence, if we think of model theory as a way to provide counterexamples to implications, then these two theories are equivalent as model theories. However, we have seen that implication-space semantics provides an account of conceptual contents as roles in implications, while truth-maker theory allows for worldly propositions that play the same implicational roles. In this sense, the interpretants of implication-space semantics are more abstract; they correspond to equivalence classes of the interpretants of truth-maker theory.

With all this in place, it is easy to show that if we restrict our implication-space models to models that are fit for a material base $\mathfrak{B}$, thus looking at $\overset{b}{\mid\!\sim}$ as we did in the previous section, then this yields a consequence relation that coincides with the truth-maker consequence for the same material base.

Proposition 84. $\Gamma \overset{b}{\mid\!\sim} \Delta$ *just in case* $\Gamma \overset{\mathfrak{B}}{\overline{\models_{TM}}} \Delta$. *(Appendix, Proposition 113)*

Let us take stock. We can use implication-space semantics to interpret the sequent calculus NMMS and truth-maker theory in parallel ways. Both theories provide us with a space of bearers of implicational roles. These are sentences in NMMS and they are worldly propositions in truth-maker theory.[11] Moreover, both theories provide us with an implication relation over sets of bearers; this relation is $\mid\!\sim$ in NMMS and $\mathbb{I}_{\overline{\diamond}}$ (defined by

$S^{\diamond}$) in truth-maker theory. So, we can abstract conceptual contents from both theories. We can give accounts of the logical connectives in all three theories, and they have exactly corresponding effects in their respective theories. That is, the operational rules of NMMS and the semantic clauses for the connectives in truth-maker theory are both equivalent to the semantic clauses of implication-space semantics. So, the three theories are equivalent as logical theories. However, implication-space semantics is, in one sense, the most abstract theory of the three theories. For it is an intrinsic semantic theory: it uses reason relations itself to interpret the items that stand in these reason relations. The codification of reason relations in the metavocabulary of implication-space semantics does not explain these reason relations in terms of normative exclusion relations among discursive acts or alethic exclusion relations among worldly states. Implication-space semantics does not specify what the bearers of implicational roles are or what the nature of exclusion relations between them is. We have thus arrived at an intrinsic characterization of the rational forms that we encountered in NMMS and in truth-maker theory.

5.3 The Monoidal Structure of Reason Relations

Let us take one more step to an even higher level of abstraction and ask what the general structure is that allows us to interpret NMMS and truth-maker theory in implication-space conceptual role semantics. Doing so is a way to ask what general class of things can fruitfully be interpreted in implication-space semantics. And given that this implication-space semantics is the intrinsic metavocabulary of reason relations, that question is equivalent to the question of what the structural features of reason relations in general are. As it turns out, mathematically speaking, the answer to this question is that reason relations in general have the structure of sets of pairs that form a commutative monoid together with a bipartition of the monoid set. In this subsection, we want to explain this answer.

A commutative monoid is a set, M, with a binary associative and commutative operation, $\bigcirc$, on M such that there is an identity element, e, in M, that is, an element such that for all $m \in \mathsf{M}$, $m \bigcirc e = m$. Now, in the case of reason relations, the elements of our monoid are pairs of sets among which reason relations hold, and we call these pairs "candidate implications." In the theories in this book, we usually think of the operation $\bigcirc$ on candidate implications as the union of the two elements of the pairs that are candidate implications, so that $\langle x, y\rangle \bigcirc \langle u, w\rangle = \langle x \cup u, y \cup w\rangle$. However, what matters for the structure of implication-space semantics is merely that the operation $\bigcirc$ that combines candidate implications is such that $\langle x, y\rangle \bigcirc \langle u, w\rangle = \langle u, w\rangle \bigcirc \langle x, y\rangle$ and $\langle x, y\rangle \bigcirc (\langle u, w\rangle \bigcirc \langle v, z\rangle) = (\langle x, y\rangle \bigcirc \langle u, w\rangle) \bigcirc \langle v, z\rangle$: commutativity and associativity. In cases in

which the items that stand in implication relations are not sets but, say, multisets (which we use when dealing with structurally noncontractive reason relations), an operation other than union must be used here, such as multiset union.[12] As long as this underlying operation is associative and commutative, the emerging operation on pairs will also be associative and commutative.

Why do we need an operation that combines candidate implications, like $\bigcirc$, in order to formulate an implication-space semantics? We need it in order to define ranges of subjunctive robustness, which are in turn what defines equivalence classes of implications and, hence, implicational roles. Why is this operation associative and commutative? Because the ordering and grouping of premises and conclusions does not matter for implications.[13]

Given the associative and commutative operation on sets of pairs, the only thing that is missing to arrive at a commutative monoid is an identity element. In the case of sets as the relata of reason relations, the identity element is the pair $\langle \emptyset, \emptyset \rangle$; for clearly $\langle x \cup \emptyset, y \cup \emptyset \rangle = \langle x, y \rangle$. In implication-space semantics, the existence of an identity element is important because it ensures that the range of subjunctive robustness of a candidate implication determines whether the candidate implication is a good implication. For it is easy to see that a candidate implication is a good implication just in case the identity element is in the candidate implication's range of subjunctive robustness. In this way, we can ensure that there cannot be two equivalent implications, that is, implications with the same range of subjunctive robustness, of which one is good and the other is not.

So far, we have seen that if we have a set of implications and the premises and conclusions of these implications are combined in a way that is associative and commutative and there is something we can combine with any premises and conclusions without changing the result, then we have a space of implications that is a commutative monoid of pairs. In order to define ranges of subjunctive robustness and, hence, implicational roles, however, we also need a bipartition that distinguishes the *good* and the *bad* implications, so that we can say that the range of subjunctive robustness of a candidate implication is the set of candidate implications that yield good implications when combined with the target candidate implication.

Let us sum up. The structure of a commutative monoid of pairs with bipartition is necessary for an implication frame. For any implication frame $\langle \mathtt{B}, \mathbb{I} \rangle$, the set $\mathbb{S} = \mathcal{P}(\mathtt{B}) \times \mathcal{P}(\mathtt{B})$ with the operation of $\langle x, y \rangle \bigcirc \langle u, w \rangle = \langle x \cup u, y \cup w \rangle$ and the identity element $\langle \emptyset, \emptyset \rangle$ is a commutative monoid and $\mathbb{I}$ is a particular subset of the monoid set. Going in the other direction, once we have a commutative monoid of pairs with bipartition, then we can define ranges of subjunctive robustness and, hence, implicational roles. For

let $\langle \mathsf{M}, \bigcirc, \mathsf{I} \rangle$ be a commutative monoid, $\langle \mathsf{M}, \bigcirc \rangle$, and a subset $\mathsf{I} \subseteq \mathsf{M}$ (which of course defines a partition on M). Then we can say that, for $M \subseteq \mathsf{M}$, the range of subjunctive robustness is $\mathtt{RSR}(M) = \{x : \forall m \in M\,(x \bigcirc m \in \mathsf{I})\}$. Moreover, we can say that for all $N, M \subseteq \mathsf{M}$, $N \approx M$ if and only if $\mathtt{RSR}(N) = \mathtt{RSR}(M)$. And then we can define implicational roles in the structure $\langle \mathsf{M}, \bigcirc, \mathsf{I} \rangle$ as above. The definitions of interpretation functions, models, and entailment then all follow as above. We accordingly arrive at an implication-space model.

So we have a structure in which we can identify implicational roles (in the way in which we do in implication-space semantics) just in case the structure is a commutative monoid of pairs with a bipartition. It is precisely in such structures that it makes sense to speak of bearers of implicational roles. Since reason relations are relations among bearers of implicational roles, this is the general structure of reason relations.

5.4 Noncontractive Reason Relations

In this section, we demonstrate the flexibility of the apparatus of implication-space semantics by extending its expressive power to capture structurally open reason relations in which Contraction fails in implication-space semantics. We discuss two examples, namely the sequent calculus $\mathsf{NMMS}^{/ctr}_{\mathfrak{B}}$ and multiplicative additive linear logic.

5.4.1 Noncontractive Material Consequence Relations

In Chapter Three, we formulated not only the sequent calculus $\mathsf{NMMS}_{\mathfrak{B}}$, which allows failures of Monotonicity and Cut, but also $\mathsf{NMMS}^{/ctr}_{\mathfrak{B}}$, in which Contraction can fail. We are now in a position to formulate an implication-space conceptual role semantics for the sequent calculus $\mathsf{NMMS}^{/ctr}_{\mathfrak{B}}$. We start by thinking of implication frames as pairs of bearers and relations between multi-sets of bearers, thus thinking of implication among bearers as a relation between multi-sets rather than sets. The definitions of ranges of subjunctive robustness, adjunction, symjunction, and conceptual contents are adjusted accordingly, always replacing set-theoretic notions applied to sets of bearers with their counterparts for multi-sets, which keep track of the number of occurrences of their elements.

Next, we change the semantic clauses for the logical vocabulary, thus tweaking our notions of interpretation functions and models. A noncontractive model is a pair of such conceptual contents and a noncontractive interpretation function, which is defined as follows:

Definition 85 (Noncontractive Interpretation Function, $[\![\cdot]\!]$)**.** An interpretation function $[\![\cdot]\!]$ is defined inductively, and maps sentences of a language

$\mathfrak{L}$ to conceptual contents. If $A \in \mathfrak{L}$ is an atomic sentence, then $[\![A]\!] =_{df.} \langle \mathtt{a}^+, \mathtt{a}^- \rangle \in \mathbb{C}$. The connective clauses are as follows:

$[\![\neg A]\!] =_{df.} \langle \mathtt{a}^-, \mathtt{a}^+ \rangle$,
$[\![A \rightarrow B]\!] =_{df.} \langle \mathtt{a}^- \sqcap \mathtt{b}^+, \mathtt{a}^+ \sqcup \mathtt{b}^- \rangle$,
$[\![A \barwedge B]\!] =_{df.} \langle \mathtt{a}^+ \sqcup \mathtt{b}^+, \mathtt{a}^- \sqcap \mathtt{b}^- \rangle$,
$[\![A \bar{\vee} B]\!] =_{df.} \langle \mathtt{a}^+ \sqcap \mathtt{b}^+, \mathtt{a}^- \sqcup \mathtt{b}^- \rangle$.

Interpretations of sets of sentences are the set of the interpretants of the sentences, for example: $\mathtt{G} = [\![\Gamma]\!] = \{[\![\gamma_i]\!] \mid \gamma_i \in \Gamma\}$.

Furthermore, we adjust our notion of model-theoretic implication as follows:

Definition 86 (Non-contractive implication in m-models, $\overset{m}{\mid\!\sim}$)**.** Given a subset of all non-contractive implication-space models m, for a given language, we say that the sentences in Γ model-theoretically imply in m the sentences in Δ, written $\Gamma \overset{m}{\mid\!\sim}_{/ctr} \Delta$, if and only if $\Gamma \overset{\mathcal{M}}{\mid\!\sim} \Delta$ for all $\mathcal{M} \in m$.

The so-defined implication-space semantics is sound and complete with respect to the sequent calculus $\mathsf{NMMS}^{/ctr}_{\mathfrak{B}}$ from Chapter Three.

Theorem 87. *Let* $\mathfrak{B} = \left\langle \mathfrak{L}_{\mathfrak{B}}, \mid\!\sim_{\mathfrak{B}} \right\rangle$ *be a base vocabulary in which contraction may fail, let* b *be the set of models that are fit for* $\mathfrak{B}$*, and let* $\Gamma, \Delta \subseteq \mathfrak{L}$*. Then* $\Gamma \overset{b}{\mid\!\sim}_{/ctr} \Delta$ *if and only if* $\Gamma \succ \Delta$ *is derivable in* $\mathsf{NMMS}^{/ctr}_{\mathfrak{B}}$*. (Appendix, Theorem 114)*

This theorem says that we have an implication-space semantics for the noncontractive reason relations codified by $\mathsf{NMMS}^{/ctr}$. It follows that the representation theorem from Chapter Three holds for the consequence relations formulated in noncontractive implication-space semantics, just as it does for the sequent-calculus version. So, for every set of candidate implications in the base, there is a single candidate implication in the logically extended vocabulary that holds if and only if all the candidate implications in the base hold.

It may be noted that the semantic clauses of the implication-space semantics for $\mathsf{NMMS}^{/ctr}$ are a bit simpler than those for NMMS. This difference corresponds exactly to the difference between the sequent rules with two top sequents in $\mathsf{NMMS}^{/ctr}$ and with three top sequents in NMMS. In the sequent rules of NMMS and in the semantic clauses for NMMS, the third element in the sequent rules and the semantic clauses, respectively, is needed to ensure that Contraction is admissible.

5.4.2 Linear Logic

Failures of contraction are perhaps most familiar from linear logic. Hence, one might wonder whether one can recover linear logic in our implication-space semantics. Indeed, one can. We will not consider the so-called exponential connectives or quantifiers here, thus limiting ourselves to *multiplicative additive linear logic* (known as "MALL"). More specifically, we will give semantic clauses for multiplicative conjunction ($\otimes$, also known as "tensor" or "times"), multiplicative disjunction ($\parr$, also known as "par"), additive conjunction (&, also known as "with"), and additive disjunction ($\oplus$, also known as "plus").

In order to formulate MALL in implication-space semantics, we leave the noncontractive models unchanged except for the semantic clauses. In particular, we change the definition of interpretation functions as follows:

Definition 88 (LL Interpretation Function, $[\![\cdot]\!]$)**.** An interpretation function $[\![\cdot]\!]$ is defined inductively, and maps sentences of a language $\mathfrak{L}$ to conceptual contents. If $A \in \mathfrak{L}$ is an atomic sentence, then $[\![A]\!] =_{df.} \langle \mathtt{a}^+, \mathtt{a}^- \rangle \in \mathbb{C}$. The connective clauses are as follows:

$[\![\neg A]\!] =_{df.} \langle \mathtt{a}^-, \mathtt{a}^+ \rangle$,
$[\![A \otimes B]\!] =_{df.} \langle \mathtt{a}^+ \sqcup \mathtt{b}^+, \mathtt{RSR}(\mathtt{RSR}(\mathtt{a}^-) \sqcup \mathtt{RSR}(\mathtt{b}^-)) \rangle$,
$[\![A \parr B]\!] =_{df.} \langle \mathtt{RSR}(\mathtt{RSR}(\mathtt{a}^+) \sqcup \mathtt{RSR}(\mathtt{b}^+)),\ \mathtt{a}^- \sqcup \mathtt{b}^- \rangle$,
$[\![A \& B]\!] =_{df.} \langle \mathtt{RSR}(\mathtt{RSR}(\mathtt{a}^+) \cup \mathtt{RSR}(\mathtt{b}^+)),\ \mathtt{a}^- \sqcap \mathtt{b}^- \rangle$,
$[\![A \oplus B]\!] =_{df.} \langle \mathtt{a}^+ \sqcap \mathtt{b}^+,\ \mathtt{RSR}(\mathtt{RSR}(\mathtt{a}^-) \cup \mathtt{RSR}(\mathtt{b}^-)) \rangle$,

Interpretations of sets of sentences are the set of the interpretants of the sentences, for example, $\mathtt{G} = [\![\Gamma]\!] = \{[\![\gamma_i]\!] \mid \gamma_i \in \Gamma\}$.

The two conjunctions of linear logic each share one of their roles with the conjunction of $\mathsf{NMMS}^{/ctr}$. The premisory role of a multiplicative conjunction $A \otimes B$ is the premisory role of the conjunction of $\mathsf{NMMS}^{/ctr}$ (and also of NMMS), namely $\mathtt{a}^+ \sqcup \mathtt{b}^+$. The conclusory role of the additive conjunction $A \& B$ is the conclusory role of the conjunction of $\mathsf{NMMS}^{/ctr}$, namely $\mathtt{a}^- \sqcap \mathtt{b}^-$. Similarly, the two disjunctions of linear logic each share one of their roles with the disjunction of $\mathsf{NMMS}^{/ctr}$. The conclusory role of the multiplicative disjunction $A \parr B$ is the conclusory role of the disjunction of $\mathsf{NMMS}^{/ctr}$ (and also of NMMS), namely $\mathtt{a}^- \sqcup \mathtt{b}^-$. The premisory role of the additive disjunction $A \oplus B$ is the premisory role of the disjunction of $\mathsf{NMMS}^{/ctr}$, namely $\mathtt{a}^+ \sqcap \mathtt{b}^+$.

The differences between MALL and $\mathsf{NMMS}^{/ctr}$ lie in the other roles, which are all the roles whose clauses (directly) mention ranges of subjunctive robustness, namely the conclusory roles of multiplicative conjunction and additive disjunction and the premisory roles of multiplicative disjunction and additive conjunction. All the implicational roles that are new in

linear logic are defined as roles of ranges of subjunctive robustness of combinations of ranges of subjunctive robustness, where this combination can happen either by adjunction or by set union. The combination happens by adjunction for the multiplicative connectives, and it happens by set union for the additive connectives.

It is noteworthy that, for each of the four linear connectives, if we allow ourselves to replace the range of subjunctive robustness of a premisory (conclusory) role with the corresponding conclusory (premisory) role and the other way around, then we can move back and forth between the premisory and the conclusory roles by taking the range of subjunctive robustness of the role with which we begin. For instance, $\mathtt{a}^+ \sqcup \mathtt{b}^+$ is the premisory role of $A \otimes B$. Replacing the premisory roles with ranges of subjunctive robustness of conclusory roles yields $\mathtt{RSR}(\mathtt{a}^-) \sqcup \mathtt{RSR}(\mathtt{b}^-)$. If we now take the role of the range of subjunctive robustness of this result, we get $\mathtt{RSR}(\mathtt{RSR}(\mathtt{a}^-) \sqcup \mathtt{RSR}(\mathtt{b}^-))$, which is the conclusory role of $A \otimes B$. The same works in the other direction, for $\mathtt{RSR}(\mathtt{RSR}(\mathtt{RSR}(\mathtt{a}^-) \sqcup \mathtt{RSR}(\mathtt{b}^-)))$ is $\mathtt{RSR}(\mathtt{a}^-) \sqcup \mathtt{RSR}(\mathtt{b}^-)$, from which we can move to $\mathtt{a}^+ \sqcup \mathtt{b}^+$. This works for all four connectives. In this way, we can compute the roles that are new in linear logic, relative to $\mathsf{NMMS}^{/ctr}$, in a uniform way.

An important upshot of this last observation is that if the sentences in linear logic are such that the premisory roles are the roles of ranges of subjunctive robustness of their conclusory roles and the other way around, then this is preserved by the linear connectives. Note, however, that our negation rule, which remains unchanged in linear logic, encodes the idea that negating something amounts to swapping the premisory and the conclusory role of what one negates. Hence, it now seems that there are two ways to switch between premisory and conclusory roles of something, namely, first, by taking its negation and, second, by taking the role of its range of subjunctive robustness. This is the crucial conceptual insight that allows us to move between the implication-space semantics for NMMS and linear logic. As will become clearer below, one way to understand implication-space semantics is to view it as what happens to phase-space semantics for linear logic if one allows negation and ranges of subjunctive robustness to come apart. The result of making this distinction is increased expressive power.

With the intepretations of multiplicative and additive conjunction and disjunction defined in this way, we restrict our space of models to those in which $\bigcup(\mathtt{a}^+ \sqcup \mathtt{a}^-) \subseteq \mathbb{I}$ for all roles of sentences $[\![A]\!] = \mathtt{a}$ and, hence A, we have $A \overset{\mathcal{M}}{\vdash\!\sim} A$, and we use "$\overset{LL}{\vdash\!\sim}$" for implications that hold in all of these models. The so-defined consequence relation is the consequence relation of MALL.

Theorem 89. *A is a linear tautology if and only if* $\overset{\mathcal{M}}{\vdash\!\sim}$ *A in all LL implication-space models. And, hence,* $\Gamma \models \Delta$ *in* MALL *if and only if* $\Gamma \overset{LL}{\vdash\!\sim} \Delta$. *(Appendix, Theorem 128)*

It might be illuminating to give some hints regarding why this theorem is true (without going through the details of the proof in the Appendix). Girard's (1987) phase-space semantics for MALL starts with a definition of a phase space, which is a commutative monoid of set $\boldsymbol{P}$, a commutative and associative operation $\bullet$ on that set, an identity element 1, and a subset $\perp_{\boldsymbol{P}} \subseteq \boldsymbol{P}$. Girard then defines the so-called dual, $G^{\perp}$, of subsets of $\boldsymbol{P}$ as follows: $G^{\perp}$ is $\{p \in \boldsymbol{P} \mid \forall q \in G\,(p \bullet q \in \perp_{\boldsymbol{P}})\}$. That is, the dual of G is the set of elements of $\boldsymbol{P}$ whose combination with anything in G is in $\perp_{\boldsymbol{P}}$. So, if we think of $\boldsymbol{P}$ as the set of candidate implications and $\perp_{\boldsymbol{P}}$ as the good implications, then the dual of something is its range of subjunctive robustness. Girard then restricts his attention to so-called facts, which are subsets of $\boldsymbol{P}$ such that $G = G^{\perp\perp}$. There is a mapping between such structures and implication frames.

- A phase space corresponds to an implication frame, where $\boldsymbol{P}$ corresponds to the space of implications $\mathbb{S}$, the anti-phases $\perp_{\boldsymbol{P}}$ correspond to the good implications $\mathbb{I}$ and the identity element, 1, of the phase space corresponds to $\mathsf{e} = \langle \emptyset, \emptyset \rangle$.
- The dual of subsets of $\boldsymbol{P}$ corresponds to the range of subjunctive robustness of subsets of $\mathbb{S}$.
- Restricting our attention to facts corresponds to working with roles of candidate implications rather than simply candidate implications.

We can thus map phase spaces to implication frames and the other way around. This correspondence between phase spaces and implication frames reveals that implication-space semantics is a close relative of phase-space semantics.[14] The proof of Theorem 89 relies on this correspondence, which is developed in more detail in the Appendix.

The most important difference between implication-space semantics and phase-space semantics is that, in Girard's phase-space semantics, the dual of facts perform two tasks at once, namely the task of providing a treatment of negation and the task of serving, in effect, as ranges of subjunctive robustness. By contrast, in implication-space semantics, these two tasks are played by distinct items. The two-sidedness of implicational roles in implication-space semantics takes over the first of these tasks, while ranges of subjunctive robustness perform the second task. If we restrict our attention to LL models in implication-space semantics, the difference between these two items is elided, which comes out in the fact that in the

canonical LL models, which are the models such that $\bigcup(\mathtt{a}^+ \sqcup \mathtt{a}^-) \subseteq \mathbb{I}$ for all roles of sentences $[\![A]\!] = \mathtt{a}$ and nothing else is in $\mathbb{I}$, we have $\mathtt{RSR}(\mathtt{a}^+) = \mathtt{a}^-$ and $\mathtt{RSR}(\mathtt{a}^-) = \mathtt{a}^+$ (Appendix: Lemma 125). Hence, in the canonical LL models $\mathtt{RSR}(\mathtt{a}) = \neg\mathtt{a}$. This does not hold in implication-space models in general, but it is built deeply into the structure of phase-space models. Thus, implication-space semantics is a generalization of Girard's phase-space semantics, namely a generalization in which ranges of subjunctive robustness and the interpretation of negation need not coincide.

This generalization allows us to include material implications from the nonlogical base consequence relation in implication-space semantics, and it allows for failures of Cut. To see this, note that if only implications whose role has the form $\mathtt{a}^+ \sqcup \mathtt{a}^-$ are good implications, as is the case in the implication-space models that codify linear logic, then the only thing in the range of subjunctive robustness of occurrences of a premisory role of some bearer are occurrences of the conclusory role of the same bearer. That is only implications whose role is that of an implication of the form $\phi \mathrel{|\!\sim} \phi$. This implies that only formally good implications are included as good implications in linear logic. Materially good implications are excluded. It is of the essence of logical expressivism to reject such an expressive restriction.

Moreover, since the instances of $\phi \mathrel{|\!\sim} \phi$ are closed under Cut, it is built into the phase structures of linear logic that Cut holds. So, the phase-space semantics for linear logic cannot capture open reason relations. If one allowed implications whose roles are not of the form $\mathtt{a}^+ \sqcup \mathtt{a}^-$ to be good implications, then the range of subjunctive robustness of bearers is no longer guaranteed to coincide with the role of their negations. Hence, the dual of a fact can no longer serve as the semantic interpretant of negation. Therefore, phase-space semantics can capture neither material nor nontransitive reason relations. It is accordingly doubly unfit to serve in the logical expressivist project of codifying open-structured reason relations. Unlike phase-space semantics, implication-space semantics allows us to give separate treatments of negation and ranges of subjunctive robustness. In implication-space semantics, we can codify, for instance, a premisory role such that the range of subjunctive robustness of its instances includes more than just instances of their conclusory role, which allows us to codify materially good implications.

As a final remark on the implication-space semantics for linear logic, it is worth pointing out that the semantic clauses for linear logic given above illustrate how we can capture sequent rules in which contexts of top sequents are combined and in which sentences that do not occur anywhere in a top sequent are introduced as subformulae in the bottom sequent. As an example of the first phenomenon, note that the conclusory role of $A \otimes B$

is $\mathtt{RSR}(\mathtt{RSR}(\mathtt{a}^-) \sqcup \mathtt{RSR}(\mathtt{b}^-))$, and that this corresponds to the following right-rule in a two-sided sequent calculus for linear logic:

$$\frac{\Gamma \succ A, \Delta \qquad \Theta \succ B, \Lambda}{\Gamma, \Theta \succ A \otimes B, \Delta, \Lambda}\ [\otimes\mathrm{R}]$$

To see how this works, it is helpful to note that $\mathtt{RSR}(\mathtt{a}^-)$ is the set of contexts $\langle \Gamma, \Delta \rangle$ such that $\Gamma \succ A, \Delta$ is a derivable sequent.[15] So $\mathtt{RSR}(\mathtt{a}^-) \sqcup \mathtt{RSR}(\mathtt{b}^-)$ is the set of contexts $\langle \Gamma \cup \Theta, \Delta \cup \Lambda \rangle$ such that $\Gamma \succ A, \Delta$ and $\Theta \succ B, \Lambda$ are both derivable. Therefore, $\mathtt{RSR}(\mathtt{RSR}(\mathtt{a}^-) \sqcup \mathtt{RSR}(\mathtt{b}^-))$ as a conclusory role is the role played by sentences, C, such that $\Gamma, \Theta \succ C, \Delta, \Lambda$ is derivable just in case $\Gamma \succ A, \Delta$ and $\Theta \succ B, \Lambda$ are both derivable.

As an example of the second phenomenon, note that the conclusory role of $A_1 \oplus B$ is $\mathtt{RSR}(\mathtt{RSR}(\mathtt{a}^-) \cup \mathtt{RSR}(\mathtt{b}^-))$, and that this corresponds to the following right-rule in a two-sided sequent calculus for linear logic:

$$\frac{\Gamma \succ A_i, \Delta}{\Gamma \succ A_1 \oplus A_2, \Delta}\ [\oplus\mathrm{R}], \mathrm{i{=}1\ or\ 2}$$

As before, $\mathtt{RSR}(\mathtt{a}^-)$ is the set of contexts $\langle \Gamma, \Delta \rangle$ such that $\Gamma \succ A, \Delta$ is a derivable sequent. And $\mathtt{RSR}(\mathtt{b}^-)$ is the set of contexts $\langle \Gamma, \Delta \rangle$ such that $\Gamma \succ B, \Delta$ is derivable. So $\mathtt{RSR}(\mathtt{a}^-) \cup \mathtt{RSR}(\mathtt{b}^-)$ is the set of contexts $\langle \Gamma, \Delta \rangle$ such that either $\Gamma \succ A, \Delta$ or $\Gamma \succ B, \Delta$ is derivable. Therefore, $\mathtt{RSR}(\mathtt{RSR}(\mathtt{a}^-) \cup \mathtt{RSR}(\mathtt{b}^-))$ as a conclusory role is the role played by sentences, C, such that $\Gamma \succ C, \Delta$ is derivable just in case $\Gamma \succ A, \Delta$ or $\Gamma \succ B, \Delta$ is derivable.

One consequence of the last point is that, unlike what happens in all variants of NMMS, if one extends a base consequence relation by adding the connectives of linear logic, then sequents that feature linear logic connectives do not always express a unique set of sequents in the base consequence relation. Rather, they sometimes express disjunctions of sequents in the base consequence relation. That is, the sequent featuring logically complex sentences is derivable just in case one of several distinct sets of nonlogical sequents are subsets of the base consequence relation. Thus, one might know that a sequent featuring logically complex sentences holds without being in a position to know that any particular sequent is in the base consequence relation. And while one can extend nontransitive base consequence relations by adding the logical vocabulary of linear logic, the Cut rule looms large in the motivation and applications of linear logic. Finally, the nonmonotonicity for which linear logic allows is best understood as the nonmonotonicity of not including superfluous premises, not the nonmonotonicity of defeasible material inferences. NMMS is accordingly vastly to be preferred to linear logic as an expressive logic.

To sum up, we can formulate an implication-space semantics for multiplicative additive linear logic. This illustrates how the familiar additive and multiplicative connectives can be formulated in implication-space semantics. However, it is built into the phase structures of linear logic that the role of the negation of a sentence and the role of its range of subjunctive robustness are the same. This implies that the phase-space semantics for linear logic can codify neither material nor non-transitive reason relations. As we have seen above, however, implication-space semantics can codify open reason relations that include material reason relations. So implication-space semantics is more flexible and can capture more kinds of reason relations—and is in this sense expressively more powerful—than the phase-space semantics for linear logic that originally inspired it.

5.5 Implicational Role Inclusion

We have noted above that two bearers of roles have the same implicational role just in case they can always be replaced for one another as premise and also as conclusion *salva consequentia*, that is, without turning a good implication into a bad one. This holds because two bearers with the same range of subjunctive robustness can always be replaced for each other, *salva consequentia*, as premises and as conclusions.

Playing the same implicational role is, however, a particularly strong version of a more general family of relations of substitutability *salva consequentia*. For it may happen that, for any X and Y, if $X, A \mid\!\sim Y$, then $X, B \mid\!\sim Y$, so that we can always replace A by B, *salva consequentia,* as a premise. And this may be true while we cannot always replace B by A or replace either of them for the other as a conclusion. Indeed, since we do not assume Cut, and in particular not a context-mixing (that is, multiplicative) version of Cut, it can happen that we can replace A by B as a premise but we cannot replace B by A as a conclusion.[16] In general, there can be different cases of substitutability *salva consequentia* that fall short of sharing a full implicational role. Since contents and implicational roles more generally are defined in terms of ranges of subjunctive robustness, we can think of such inclusions among ranges of subjunctive robustness as a kind of inclusion relation among contents and other implicational roles (although not in an immediate set-theoretic sense). In this section, we want to investigate phenomena like that of an implicational role including another implicational role. In this way, we study relations of substitutability *salva consequentia* in more detail and in greater generality. It turns out that this allows us to connect implication-space semantics to several familiar logical theories.

5.5.1 Foundations of Role Inclusion

In order to study the phenomenon of substitutability *salva consequentia* of bearers of roles, it proves useful to define a notion of role inclusion.[17] The idea is that the role $\mathcal{R}(\alpha)$ is included in the role $\mathcal{R}(\beta)$ if and only if we can always replace, *salva consequentia*, something that plays the first role with something that plays the second role. This is the case if and only if the range of subjunctive robustness of anything in $\mathcal{R}(\alpha)$ is a subset of the range of subjunctive robustness of anything in $\mathcal{R}(\beta)$. To see this, take the premisory role of a bearer A: $\mathcal{R}(\langle\{A\},\emptyset\rangle)$. Its range of subjunctive robustness consists of all sets of candidate implications such that for all their elements, $\langle X,Y\rangle$, we have $\langle X\cup\{A\},Y\rangle \in \mathbb{I}$. Now, suppose that all of these sets of candidate implications are also in the range of subjunctive robustness of the premisory role of some other bearer B. That is, they are in the range of subjunctive robustness of $\mathcal{R}(\langle\{B\},\emptyset\rangle)$. Then for every candidate implication, $\langle X,Y\rangle$, if $\langle X\cup\{A\},Y\rangle \in \mathbb{I}$, then $\langle X\cup\{B\},Y\rangle \in \mathbb{I}$, since $\langle X\cup\{A\},Y\rangle \in \mathbb{I}$ holds just in case$\{\langle X,Y\rangle\}$ is in the range of subjunctive robustness of $\mathcal{R}(\langle\{A\},\emptyset\rangle)$ and analogously for $\langle X\cup\{B\},Y\rangle \in \mathbb{I}$. In other words, for any X and Y, if $X,A \mid\!\sim Y$, then $X,B \mid\!\sim Y$. Thus, we can always replace, *salva consequentia*, A as a premise with B as a premise. In such a case we say that the premisory role of B is included in the premisory role of A. In general, we say that a role is included in another role just in case the range of subjunctive robustness of the first role is a subset of the range of subjunctive robustness of the second role, in all models under consideration.

It will prove useful to generalize this idea to cases in which several roles are included in several other roles. The following way of doing this proves fruitful: Take for example the premisory role of A and the conclusory role of B. They are included in the combined conclusory roles of C and D, in a class of models, if and only if we have $\langle X,Y\cup\{C,D\}\rangle \in \mathbb{I}$ whenever we have $\langle X\cup\{A\},Y\rangle \in \mathbb{I}$ and $\langle X,Y\cup\{B\}\rangle \in \mathbb{I}$. To put it differently, this role inclusion holds just in case everything that is in the range of subjunctive robustness of all the included roles is also in the range of subjunctive robustness of the adjunction of the including roles. We thus take the symjunction of the included roles and the adjunction of the including roles. This particular way for handling sets of implicational roles makes the relations below easiest to see, and it can helpfully be related to proofs in sequent calculi. The following definition makes our notion of implicational role inclusion precise:

Definition 90 (Implicational role inclusion, $\preceq$)**.** Given a set of implication-space models, m, we say that the set of implicational roles $\{\mathcal{R}_1,...,\mathcal{R}_n\}$ is included in the set of implicational roles $\{\mathcal{R}_k,...,\mathcal{R}_l\}$ in m, written

$\mathcal{R}_1, ..., \mathcal{R}_n \preceq \mathcal{R}_k, ..., \mathcal{R}_l$ in m, just in case $\mathtt{RSR}(\bigsqcap_{i=1}^{n} \mathcal{R}_i) \subseteq \mathtt{RSR}(\bigsqcup_{j=k}^{l} \mathcal{R}_j)$ in every model in m.

Thus, everything that makes the symjunction of the included roles good also makes the adjunction of the including roles good. To see what this amounts to, let us start by noting some general features of implicational role inclusion. In particular, we can note how structural principles governing implications of the underlying implication space show up in role inclusions among contents. The following role inclusions hold, in a model, if and only if the implications among the bearers with these contents obey Containment, Monotonicity, and Cut, respectively[18]:

Containment: $\forall \mathtt{p} \in \mathbb{C}\ (\star \preceq \mathtt{p}^+, \mathtt{p}^-)$.

Monotonicity: $\forall \mathtt{p}, \mathtt{q} \in \mathbb{C}\ (\mathtt{p}^{+/-} \preceq \mathtt{p}^{+/-}, \mathtt{q}^{+/-})$

Cut: $\forall \mathtt{p} \in \mathbb{C}\ (\mathtt{p}^+, \mathtt{p}^- \preceq \mathtt{e})$

To see why these are formulations of the respective structural principles, recall that $\mathtt{e} = \langle \emptyset, \emptyset \rangle$ is the empty candidate implication, and $\star = \langle \mathtt{B}, \mathtt{B} \rangle$ is the implication whose premise-set and conclusion-set is the set of all bearers of the implication frame. So the range of subjunctive robustness of $\mathtt{e}$ is the set of good implications $\mathbb{I}$. And the range of subjunctive robustness of $\star$ is the entire implication space $\mathbb{S}$. After all, $\langle \mathtt{B}, \mathtt{B} \rangle \in \mathbb{I}$ and $\langle \mathtt{B}, \mathtt{B} \rangle$ does not change by taking the set-union of its premises or conclusions with any set of bearers. So our formulation of Containment says that whenever we combine the premisory and the conclusory role of any content of a bearer in one implication (by taking their adjunction), the result has the maximal range of subjunctive robustness. That is, all implications in which any bearers with the same content occur on the left side and also on the right side are indefeasible implications. Similarly, our formulation of Monotonicity says that anything that makes good all implications in some (premisory or conclusory) role $\mathtt{p}^{+/-}$ also makes good all implications, $\mathtt{p}^{+/-} \sqcup \mathtt{q}^{+/-}$, with additional premises and conclusions. Finally, our formulation of Cut says that if an implication is made good by adjoining the premisory role, $\mathtt{p}^+$, of a content and it is also made good by adjoining the conclusory role, $\mathtt{p}^-$, of that content, then the implication has a range of subjunctive robustness that is included in the range of subjunctive robustness of $\mathtt{e}$. But since $\mathtt{RSR}(\mathtt{e}) = \mathbb{I}$, this just means that the implication is already good by itself.

We can now see that, in effect, all of our semantic clauses are statements of role identities, if we extend the clauses to cover "logically complex" roles that might not have bearers. They translate immediately into mutual

role inclusions, namely as follows, where "$\preceq\succeq$" stands for a role inclusion relation that goes in both directions:

Fact 91. *By our semantic clauses for the implication-space semantics of* NMMS *(Definition 70),*[19] *for any contents* $\mathtt{a} = \langle \mathtt{a}^+, \mathtt{a}^- \rangle$ *and* $\mathtt{b} = \langle \mathtt{b}^+, \mathtt{b}^- \rangle$, *we have* $\mathtt{a}^+ \preceq\succeq (\neg\mathtt{a})^-$, *and* $\mathtt{a}^- \preceq\succeq (\neg\mathtt{a})^+$, *and* $\mathtt{a}^+ \sqcup \mathtt{b}^+ \preceq\succeq (\mathtt{a} \wedge \mathtt{b})^+$, *and* $\mathtt{a}^- \sqcap \mathtt{b}^- \sqcap (\mathtt{a}^- \sqcup \mathtt{b}^-) \preceq\succeq (\mathtt{a} \wedge \mathtt{b})^-$, *and similarly for the other semantic clauses.*

For example, our semantic clauses say $[\![\neg A]\!] =_{df.} \langle \mathtt{a}^-, \mathtt{a}^+ \rangle$, so that $(\neg\mathtt{a})^- = \mathtt{a}^+$ and $(\neg\mathtt{a})^+ = \mathtt{a}^-$. The resulting mutual role inclusions are $\mathtt{a}^+ \preceq\succeq (\neg\mathtt{a})^-$ and $\mathtt{a}^- \preceq\succeq (\neg\mathtt{a})^+$. And the other semantic clauses imply mutual role inclusions in an analogous way. It may also be helpful to note the following immediate consequences of the definition of implicational role inclusion:

Fact 92. *For any roles* $\mathcal{R}\langle \Gamma, \Delta \rangle$ *and* $\mathcal{R}\langle \Theta, \Lambda \rangle$*:*

(i) $\mathcal{R}\langle \Gamma, \Delta \rangle \preceq \mathtt{e}$ *just in case* $\mathtt{RSR}\langle \Gamma, \Delta \rangle \subseteq \mathbb{I}$.

(ii) $\star \preceq \mathcal{R}\langle \Gamma, \Delta \rangle$ *just in case, for all implications* $\langle X, Y \rangle$, $\langle \Gamma \cup X, \Delta \cup Y \rangle \in \mathbb{I}$, *that is,* $\langle \Gamma, \Delta \rangle \in \mathbb{I}$ *holds persistently.*

(iii) $\mathcal{R}\langle \Gamma, \Delta \rangle \preceq \mathtt{e}, \mathcal{R}\langle \Theta, \Lambda \rangle$ *just in case* $\mathcal{R}\langle \Gamma, \Delta \rangle \preceq \mathcal{R}\langle \Theta, \Lambda \rangle$.

(iv) $\star, \mathcal{R}\langle \Gamma, \Delta \rangle \preceq \mathcal{R}\langle \Theta, \Lambda \rangle$ *just in case* $\mathcal{R}\langle \Gamma, \Delta \rangle \preceq \mathcal{R}\langle \Theta, \Lambda \rangle$.

Finally, we can use the notion of role inclusion just introduced to define a notion of content inclusion. Let us start with an intuitive case. There seems to be a sense in which the content of "Fido is a dog" includes the content of "Fido is a mammal." One way to spell out what this means is the following: Everything that follows from "Fido is a mammal" (together with any further premises) also follows from "Fido is a dog" (together with those same further premises). And, conversely, everything that implies "Fido is a dog" also implies "Fido is a mammal." We can formulate this idea in terms of substitution *salva consequentia* as follows: We can always substitute *salva consequentia* "Fido is a mammal" by "Fido is a dog" as premises. And we can always substitute *salva consequentia* "Fido is a dog" by "Fido is a mammal" as conclusions. And putting this in terms of role inclusion, we can say that the premisory role of "Fido is a mammal" is included in the premisory role of "Fido is a dog," and the conclusory role of "Fido is a dog" is included in the conclusory role of "Fido is a mammal." Let us formulate this idea as a general definition:

Definition 93 (Content inclusion, $\Subset$). For any contents $\mathtt{a} = \langle \mathtt{a}^+, \mathtt{a}^- \rangle \in \mathbb{C}$ and $\mathtt{b} = \langle \mathtt{b}^+, \mathtt{b}^- \rangle \in \mathbb{C}$, we say that content $\mathtt{a}$ includes content $\mathtt{b}$, written $\mathtt{b} \Subset \mathtt{a}$, if and only if $\mathtt{b}^+ \preceq \mathtt{a}^+$ and $\mathtt{a}^- \preceq \mathtt{b}^-$.

With this notion of content inclusion in hand, we can say that two contents overlap just in case there is some third content that is included in both contents.[20]

In the remainder of this section, we show that many familiar logics can be understood as logics of implicational role inclusion. We can thus see these logics as allowing us to codify different aspects of inclusions among implicational roles and contents.

5.5.2 Logics Based on the Strong Kleene Scheme

In this subsection, we show how familiar logics that use the strong Kleene truth-tables emerge as implicational role inclusions within implication-space semantics (of NMMS). These logics are Priest's paraconsistent logic of paradox (LP), the paracomplete strong Kleene logic (K3), the nontransitive strict-tolerant logic (ST), and the nonreflexive tolerant-strict logic (TS).

The logics LP, K3, ST, and TS are usually formulated by using three truth-values, often denoted by 1 (for truth), 0 (for falsehood), and $\frac{1}{2}$ for the third truth-value. All four logics are based on the strong Kleene truth-tables, which are the following (treating the conditional as defined): the value of a negation $v(\neg\phi)$ is $1 - v(\phi)$; the value of a conjunction is the minimum of the values among the conjuncts, and the value of a disjunction is the maximum of the values among the disjuncts. An interpretation is a function that assigns to each sentence of a language exactly one of these three truth-values, in a way that respects the strong Kleene truth-tables.

The only difference between LP, K3, ST, and TS is the way in which consequence is defined in them. If we define consequence as preservation of values more than 0, this yields the logic LP (Priest, 2006). And if we define consequence as preservation of value 1, then this yields the logic K3. A consequence relation holds in ST just in case every interpretation in which all the premises have value 1 is such that some conclusion has a value other than 0 (Ripley, 2012; Cobreros et al., 2012, 2020a). And a TS consequence holds just in case every interpretation in which all the premises have values other than 0 is such that some conclusion has value 1 (French, 2016; Cobreros et al., 2020a).

Famously, Priest advocates LP as a response to paradoxes, while Kripke's solution is based on K3.[21] Advocates of LP think of the third truth-value as representing truth-value-gluts, that is, as cases in which a sentence is both true and false (in a model). And advocates of K3 think of the third truth-value as representing truth-value-gaps, that is, as cases in which a sentence

is neither true nor false (in a model). If we include a paradoxical sentence, λ, that receives the third truth-value in all models, then LP invalidates the principle *ex contradictione quodlibet*, so that $\lambda \wedge \neg\lambda \not\models_{LP}$. And K3 invalidates the principle of the excluded middle, so that $\not\models_{K3} \lambda \vee \neg\lambda$.

Unless we add paradoxical sentences or the like to our language, the consequence relation of ST coincides with classical logic, and the consequence relation of TS is empty. However, one can add a transparent truth-predicate to all of these logics without making their consequence relations trivial. If one adds a paradoxical sentence, then Cut fails in ST. One can think of TS as rejecting Reflexivity (French, 2016).

It is well known that the logics LP and K3 coincide, each in its own way, with the so-called "local metainferential" consequence relation of ST (Cobreros et al., 2020a, b; Barrio et al., 2015; Dicher and Paoli, 2019). Metainferential validity is a relation between sets of candidate implications. A metainference is locally valid just in case every model that is a counterexample to all the conclusion candidate implications is a counterexample to at least one premise candidate implication. We can recover the relation of local metainferential validity in implication-space semantics. This works because there is an ST-counterexample-preserving bijection between strong Kleene valuations and a particular subset of implication-space models, which we call the conic implication-space models.

Definition 94 (Conic implication-space models)**.** A conic implication-space model, $\mathcal{M}$, is a reflexive implication space model in which there is a pair $\langle \Gamma, \Delta \rangle$ such that $X \not\models^{\mathcal{M}} Y$ if and only if $\langle X, Y \rangle \in \mathcal{P}(\Gamma) \times \mathcal{P}(\Delta)$.

In a conic implication-space model, for some $\langle \Gamma, \Delta \rangle$, every bad candidate implication is such that its premise-set is a subset of Γ and its conclusion-set is a subset of Δ. Thus, the bad candidate implications form, as it were, a cone with $\langle \Gamma, \Delta \rangle$ as a top element and all other bad candidate implications sitting in a set-theoretic inclusion lattice below that top element (where the lattice structure is such that both sets of a pair must be subsets of the elements of the pairs above it). We can map such models to strong Kleene valuations, v, in an ST-counterexample-preserving bijection as follows: $v(A) = 1$ just in case $A \in \Gamma$, $v(A) = 0$ just in case $A \in \Delta$, and $v(A) = \frac{1}{2}$ otherwise. Intuitively, the bad implications are those with true premises and false conclusions, and $\langle \Gamma, \Delta \rangle$, at the top of the cone, consists of the pair of *all* the true premises and *all* the false conclusions. So every bad implication has premises and conclusions that are subsets of these and so are found lower in the cone on the set-inclusion lattice. It is easy to see that this bijection maps implications for which the strong Kleene model is an ST counterexample to implications that are bad in the conic implication-space

model, and the other way around. For every conic implication-space model, we can construct the strong Kleene model to which it is mapped by our bijection, and the other way around. And since the consequence relation of ST is reflexive, the consequence relations of these conic implication-space models also obey Reflexivity (see Appendix, Theorem 135).

The well-known connections between ST, LP, K3, and TS emerge in implication-space semantics in terms of implicational role inclusions over conic models (see Appendix, Fact 137). In particular, LP agrees with role inclusions among conclusory roles, and K3 agrees with the converse of role inclusions among premisory roles, over all conic implication-space models. ST coincides with the role of implications that include the role of the whole implication space. And TS coincides with the inclusion of the symjunction of the roles of the conjunction of the premises and the disjunction of the conclusions in the role of the identity element.

Proposition 95. *Suppose that the only constraint on implication-space models is that they be conic, and let* $[\![\Gamma]\!] = \mathtt{G}$ *and* $\mathtt{G}^+ = \{\mathtt{g}^+ \mid \mathtt{g} \in \mathtt{G}\}$ *and* $\mathtt{G}^- = \{\mathtt{g}^- \mid \mathtt{g} \in \mathtt{G}\}$ *and analogously for* $[\![\Delta]\!] = \mathtt{D}$. *Then:*

$\Gamma \models_{LP} \Delta$ *if and only if* $(\bigwedge \mathtt{G})^- \preceq \mathtt{D}^-$.

$\Gamma \models_{K3} \Delta$ *if and only if* $(\bigvee \mathtt{D})^+ \preceq \mathtt{G}^+$.

$\Gamma \models_{ST} \Delta$ *if and only if* $\star \preceq \mathtt{G}^+, \mathtt{D}^-$.

$\Gamma \models_{TS} \Delta$ *if and only if* $(\bigwedge \mathtt{G})^-, (\bigvee \mathtt{D})^+ \preceq \mathtt{e}$. *(Appendix, Proposition 139)*

This means that LP is the logic of conclusory role inclusions. Thus, LP tells us which conclusions we can always replace *salva consequentia* for which other conclusions, in all conic implication-space models. Similarly, K3 tells us by which premises we can always replace *salva consequentia* which other premises, in all conic implication-space models. Since $\star \preceq \mathtt{G}^+, \mathtt{D}^-$ just in case $\Gamma \overset{\mathcal{M}}{\mid\!\sim} \Delta$ holds persistently in all conic models, ST tells us which implications are persistently good in all conic models. Moreover, $(\bigwedge \mathtt{G})^-, (\bigvee \mathtt{D})^+ \preceq \mathtt{e}$ holds just in case $\mathtt{RSR}((\bigwedge \mathtt{G})^- \sqcap (\bigvee \mathtt{D})^+) \subseteq \mathbb{I}$. So, a consequence $\Gamma \models_{TS} \Delta$ holds just in case, for all $g \subseteq \Gamma$, if $X \overset{\mathcal{M}}{\mid\!\sim} g, Y$, then $X \overset{\mathcal{M}}{\mid\!\sim} Y$, and for all $d \subseteq \Delta$, if $X, d \overset{\mathcal{M}}{\mid\!\sim} Y$, then $X \overset{\mathcal{M}}{\mid\!\sim} Y$. Thus, TS captures which sentences can be dropped from an implication *salva consequentia*.

Since the logics LP and K3 have received considerable attention in the literature, it may be worth commenting a bit more on how they emerge in implication-space semantics. Let us start by pointing out that, from the perspective of implication-space semantics, LP and K3 are limiting cases that look only at formal implications: implications that hold in all conic models. For a candidate implication has a counterexample in a conic model if and only if it has a counterexample in the set of models that obey Containment. Thus, we lose all information about material implication and

incompatibility. If we looked at fewer models, for instance, by requiring that certain material implication and incompatibility relations hold in the implications spaces of these models, we would get more substantive role inclusion relations. Thus, LP and K3 provide the lower bounds of conclusory and premisory role inclusions, respectively. It follows from this that, given Containment, the intersection of LP and K3, that is, the set of implications that hold in both logics, is a lower bound on content inclusion.

Proposition 96. *The content inclusion* $\mathtt{b} \Subset \mathtt{a}$ *holds in all conic models, if and only if, for all A such that* $[\![A]\!] = \mathtt{a}$ *and for all B such that* $[\![B]\!] = \mathtt{b}$*, we have* $A \models_{K3} B$ *and* $A \models_{LP} B$*.*

Let us now turn to the paraconsistent and paracomplete natures of LP and K3, respectively. The principles that LP and K3 reject, but that hold in classical logic, are related in intimate ways to substitution *salva consequentia*. In classical logic, we have $(\mathtt{a} \wedge \neg\mathtt{a})^- \preceq \mathtt{e}$ and $(\mathtt{a} \vee \neg\mathtt{a})^+ \preceq \mathtt{e}$. The first of these role inclusions can be viewed as a version of *ex contradictione quodlibet,* and it fails in LP. In particular, the role inclusion $(\mathtt{a} \wedge \neg\mathtt{a})^- \preceq \mathtt{e}$ says that, in all conic models, for all Γ and all Δ, if $\Gamma \mathrel{\overset{\mathcal{M}}{\mid\!\sim}} A \wedge \neg A, \Delta$, then $\Gamma \mathrel{\overset{\mathcal{M}}{\mid\!\sim}} \Delta$. The second can be viewed as a version of excluded middle, and it fails in K3. In particular, the role inclusion $(\mathtt{a} \wedge \neg\mathtt{a})^+ \preceq \mathtt{e}$ says that, in all conic models, for all Γ and all Δ, if $\Gamma, A \vee \neg A \mathrel{\overset{\mathcal{M}}{\mid\!\sim}} \Delta$, then $\Gamma \mathrel{\overset{\mathcal{M}}{\mid\!\sim}} \Delta$.

Given Containment, these two role inclusions are both versions of Cut in the underlying implication space. For, $(\mathtt{a} \wedge \neg\mathtt{a})^- \preceq \mathtt{e}$ if and only if $\mathtt{a}^- \sqcap \mathtt{a}^+ \sqcap (\mathtt{a}^- \sqcup \mathtt{a}^+) \preceq \mathtt{e}$, which holds just in case $\mathtt{a}^-, \mathtt{a}^+ \preceq \mathtt{e}$ and $\mathtt{a}^- \sqcup \mathtt{a}^+ \preceq \mathtt{e}$ both hold. And $\mathtt{a}^- \sqcup \mathtt{a}^+ \preceq \mathtt{e}$ is guaranteed by Containment, which holds in all conic models. Hence, $\mathtt{a}^-, \mathtt{a}^+ \preceq \mathtt{e}$ if and only if $(\mathtt{a} \wedge \neg\mathtt{a})^- \preceq \mathtt{e}$. And as we have seen above, $\mathtt{a}^-, \mathtt{a}^+ \preceq \mathtt{e}$ means that we can always apply Cut to bearers of the content $\mathtt{a}$. Similarly, supposing Containment, Cut holds in the underlying implication relation just in case the role inclusion $(\mathtt{a} \vee \neg\mathtt{a})^+ \preceq \mathtt{e}$ holds. For, $(\mathtt{a} \vee \neg\mathtt{a})^+ \preceq \mathtt{e}$ if and only if $\mathtt{a}^+ \sqcap \mathtt{a}^- \sqcap (\mathtt{a}^+ \sqcup \mathtt{a}^-) \preceq \mathtt{e}$. By the previous reasoning, it follows that $(\mathtt{a} \vee \neg\mathtt{a})^+ \preceq \mathtt{e}$ means that we can always apply Cut to bearers of the content $\mathtt{a}$.

It might seem surprising that what looks like a truth-value glut in LP shows up as a failure of Cut in implication-space semantics, and similarly for what looks like a truth-value gap in K3. Why is that? The technical answer is this: Given our connectives and Containment, the roles $(\mathtt{a} \vee \neg\mathtt{a})^+$ and $(\mathtt{a} \wedge \neg\mathtt{a})^-$ and $\mathtt{a}^- \sqcap \mathtt{a}^+$ are all equivalent. This means that, with $[\![A]\!] = \mathtt{a}$, we can always substitute, *salva consequentia*, $A \vee \neg A$ as a premise in an implication for $A \wedge \neg A$ as a conclusion, and the other way around. And we can replace either one of them, *salva consequentia*, by

two implications, which are otherwise the same, with A being a premise in one implication and A being a conclusion in the other implication. So, $\Gamma, A \vee \neg A \overset{\mathcal{M}}{\mid\!\sim} \Delta$ just in case $\Gamma \overset{\mathcal{M}}{\mid\!\sim} A \wedge \neg A, \Delta$ just in case $\Gamma \overset{\mathcal{M}}{\mid\!\sim} A, \Delta$ and $\Gamma, A \overset{\mathcal{M}}{\mid\!\sim} \Delta$ both hold. Hence, we can drop $A \vee \neg A$ as a premise or, equivalently, $A \wedge \neg A$ as a conclusion in an implication, *salva consequentia*, just in case we can apply Cut to two implications that are like the one with which we started except that A figures as a premise in one and as a conclusion in the other.

When we interpret this situation in terms of normative bilateralism, we can say that there is a sentence, A, that can neither be accepted nor rejected, in a given position, so that Cut fails, just in case, in that position, one can neither reject $A \wedge \neg A$ nor accept $A \vee \neg A$. But if we cannot reject $A \wedge \neg A$, then it can seem that we must accept A and also accept $\neg A$. And this seems to commit us to saying that A is true and also false. And if we cannot accept $A \vee \neg A$, then it seems that we must reject A and also reject $\neg A$. And this seems to commit us to saying that A is neither true nor false.

When we interpret the situation in terms of truth-maker theory, we can say that there is a sentence, A, such that neither one of its truth-makers nor one of its falsity-makers can be fused with a given state into a state that is possible, just in case neither a falsity-maker of $A \wedge \neg A$ nor a truth-maker of $A \vee \neg A$ can be fused with the given state into a state that is possible. But if $A \wedge \neg A$ cannot be made false, then it can seem that neither A nor $\neg A$ can be made false, so that it seems that A and $\neg A$ must both be made true. And this seems to commit us to saying that A is true and also false. And if $A \vee \neg A$ cannot be made true, then it seems that neither A nor $\neg A$ can be made true. And this seems to commit us to saying that A is neither true nor false.

Returning to the high level of abstraction of this chapter, the general interpretation is that if $(\mathtt{a} \wedge \neg\mathtt{a})^-$ is ruled out, then $\mathtt{a}^-$ and $(\neg\mathtt{a})^-$ are both ruled out. Thus, $\mathtt{a}^+$ and $(\neg\mathtt{a})^+$ seem forced upon us, which yields the impression of a truth-value glut. And if $(\mathtt{a} \vee \neg\mathtt{a})^+$ is ruled out, then $\mathtt{a}^+$ and $(\neg\mathtt{a})^+$ are both ruled out. And this yields the impression of a truth-value gap. Given the role equivalences above, we know that these situations arise just in case the role $\mathtt{a}^- \sqcap \mathtt{a}^+$ is ruled out, which shows up as a failure of Cut.

To sum up, the impression that there are truth-value gluts or truth-value gaps can be understood as arising from failures of Cut, when we think about implicational role inclusions as relations that preserve some designated truth-value. From the perspective of implication-space semantics, however, this impression of truth-value gluts or gaps is merely a reflection of the fact that LP codifies the relation of conclusory role

inclusion in all conic models, and K3 codifies the converse relation of premisory role inclusion in all conic models.

A comment on our recovery of ST and TS may also be in order. We can see from the proposition above that ST and TS are not dual logics, like LP and K3, which result from each other by flipping the sides of the role inclusions and switching premisory to conclusory roles, and the other way around, for the whole sets of premises and conclusions combined by adjunction. However, ST and TS would be dual logics if the premisory role of e was the conclusory role of $\star$ and the other way around. Now, the premisory role of e is not the conclusory role of $\star$. Rather, e and $\star$ are the minimal and the maximal element of the set-theoretic inclusion lattice of subsets of $\mathbb{S}$.[22] This suggests that we can think of the relation between ST and TS in terms of the inversion of the ordering of a lattice. However, we will not develop this idea here.

For our purposes, it suffices to note that, in implication-space semantics, the consequence relations of ST and TS emerge as relations of implicational role inclusion in conic models. The general fact that underlies this connection, as well as the connection between implication-space semantics and LP and K3, is that there is an ST-counterexample-preserving bijection between strong Kleene interpretations and conic implication-space models.

5.5.3 Logics of Content and Metaphysics

In this subsection, we want to add a final point to our demonstration of the power of implication-space semantics to recover and connect extant logical theories. There is a family of logics that have been considered as logics of content or logics of factual equivalence (Fine, 2016; Correia, 2016).[23] We now show how some of these logics are related to identity and inclusions of contents considered as implicational roles.

We begin by looking at a logic that Correia (2016) suggested as a "logic of factual equivalence." Correia is interested in determining when two sentences, A and B, describe the same facts or states in virtue of their (propositional) logical form, written $A \approx B$. Correia offers a Hilbert-style axiomatic system for such statements, and we use "Correia's logic" for this logic. Correia's is a proper fragment of Angell's (1989; 1977) logic of analytic entailment, which Fine (2016) advocates as a logic of content. Correia's logic differs from Angell's logic in that it doesn't validate the distributive principle according to which $A \vee (B \wedge C)$ is equivalent to $(A \vee B) \wedge (A \vee C)$. However, Correia's logic does include distribution of conjunction over disjunction as an axiom (using Correia's labels).

A10 $\quad A \wedge (B \vee C) \approx (A \wedge B) \vee (A \wedge C)$

We get Angell's first-degree system for analytic equivalence by adding the distribution principle that Correia rejects as another axiom, namely:

A11 $\quad A \vee (B \wedge C) \approx (A \vee B) \wedge (A \vee C)$

If we add A11 but not A10, we get what Correia calls his "dual" logic. We write $\vdash_c A \approx B$ to say that $A \approx B$ is a theorem of Correia's logic, and $\vdash_{dc} A \approx B$ to say that $A \approx B$ is a theorem of Correia's dual logic. We can recover both Correia's logic and his dual logic in terms of implicational role inclusions, indeed in terms of role identities, thereby giving a precise sense to their claims to be logics of content.

Proposition 97. *If the only constraint on models is that they obey the semantic clauses for the implication-space semantics of* NMMS *(Definition 70), then Correia's logic and his dual system are related to implicational roles as follows. For all* $[\![A]\!] = \mathtt{a}$ *and* $[\![B]\!] = \mathtt{b}$*:*

$\vdash_c A \approx B$ *if and only if, in all models,* $\mathtt{a}^+ \preceq\succeq \mathtt{b}^+$*, and so* $\mathtt{a}^+ = \mathtt{b}^+$*.*

$\vdash_{dc} A \approx B$ *if and only if, in all models,* $\mathtt{a}^- \preceq\succeq \mathtt{b}^-$*, and so* $\mathtt{a}^- = \mathtt{b}^-$*. (Appendix, Proposition 143)*

We can also express this in terms of content identity.

Proposition 98. *The contents* $[\![A]\!] = \mathtt{a}$ *and* $[\![B]\!] = \mathtt{b}$ *are identical, and so* $\mathtt{a} \Subset \mathtt{b}$*, in all implication-space models, if and only if* $\vdash_c A \approx B$ *and* $\vdash_{dc} A \approx B$*.*

So while mutual implication in LP and also K3 ensures that two bearers have the same content in all conic models, a more demanding connection is needed when we include nonconic models. And this more demanding connection is that the two bearers must be equivalent in Correia's logic and Correia's dual logic. In other words, we get Correia's logic from K3 as a logic of premise substitution *salva consequentia* if, firstly, we require that substitutions must be possible in both directions and, secondly, we disallow substitutions that might not be *salva consequentia* in nonconic models. And we get Correia's dual logic from LP in the same way.

Without entering into any debate on these issues, we can note that role inclusions can be used to recover some notions that are currently popular in so-called "analytic metaphysics." Correia and Skiles (2019) offer, for instance, accounts of gounding, essence, and generalized identity on the basis of Correia's logic. When we take the versions of their notions that apply at the sentential level, we can translate them into implicational role inclusions or identities. According to Correia and Skiles a generalized identity claim of the form "For it to be the case that A is for it to be the

case that B" holds in virtue of logical form just in case $\vdash_c A \approx B$. Hence, this notion coincides with the identity of premisory roles of the bearers A and B. Regarding essence, Correia and Skiles's notion of full factual essence turns out to coincide also with identity of premisory implicational roles, as they hold that a generalized identity claim connecting two sentences holds just in case the facts stated by those sentences are each other's full essence. It turns out that a fact is part of the essence of another fact, in their terms, just in case some adjunction of the premisory role of the sentence that states the first fact with another role is identical to the premisory role of a sentence that states the second fact.

Fact 99. *Factual Essence: Its being the case that A is what it is for it to be the case that B in full, in virtue of truth-functional logical form (Correia and Skiles, 2019, 651), if and only if, for* $[\![A]\!] = \mathtt{a}$ *and* $[\![B]\!] = \mathtt{b}$, *we have* $\mathtt{a}^+ = \mathtt{b}^+$ *in all implication-space models. And its being the case that A is in part what it is for it to be the case that B, in virtue of truth-functional logical form (Correia and Skiles, 2019, 651), if and only if , for* $[\![A]\!] = \mathtt{a}$ *and* $[\![B]\!] = \mathtt{b}$, *there is some* $\mathtt{c}^+$ *such that* $\mathtt{a}^+ \sqcup \mathtt{c}^+ = \mathtt{b}^+$ *in all implication-space models. (Appendix, Fact 144)*

If we wanted to include more than just those factual essence claims that hold in virtue of logical form, we could restrict ourselves to a subset of all implication-space models. Requiring models to respect the corresponding identities of premisory roles will then ensure that the desired claims about essences are true in all these models. We can translate the notion of grounding in a similar way.

Fact 100. *Strict Full Grounding: Its being the case that* $A_1, ..., A_n$ *makes it the case that B, in virtue of truth-functional logical form (Correia and Skiles, 2019, 655), if and only if, in all implication-space models, for* $[\![A_i]\!] = \mathtt{a}_i$ *and* $[\![B]\!] = \mathtt{b}$, *(i) for some* $[\![C]\!] = \mathtt{c}$, *we have* $(\mathtt{a}_1^+ \sqcup ... \sqcup \mathtt{a}_n^+) \sqcap \mathtt{c}^+ \sqcap (\mathtt{a}_1^+ \sqcup ... \sqcup \mathtt{a}_n^+ \sqcup \mathtt{c}^+) \preceq\succeq \mathtt{b}^+$ *and (ii)* $\forall 1 \leq i \leq n$ *there are no* $[\![D]\!] = \mathtt{d}$ *and* $[\![E]\!] = \mathtt{e}$ *such that* $(\mathtt{b}^+ \sqcup \mathtt{d}^+) \sqcap \mathtt{e}^+ \sqcap (\mathtt{b}^+ \sqcup \mathtt{d}^+ \sqcup \mathtt{e}^+) \preceq\succeq \mathtt{a}_i^+$. *(Appendix, Fact 145)*

Recently, Elgin (2021) has complained that Correia and Skiles's notion of generalized identity, in effect, only considers truth-makers and does not take falsity-makers into account. Elgin suggests that a generalized identity statement connecting two sentences is true just in case the two sentences have the same truth-makers and also the same falsity-makers. Elgin offers an axiomatic system for generalized identity statements that hold merely in virtue of logical form.[24] It is easy to see that Elgin's notion can be translated into implication-space semantics as follows:

Fact 101. *A sentence of the form "For it to be the case that A is for it to be the case that B" is true in virtue of logical form, in the sense of Elgin (2021, sec 4), if and only if, for* $[\![A]\!] = \mathtt{a}$ *and* $[\![B]\!] = \mathtt{b}$*, we have* $\mathtt{a} = \mathtt{b}$*, and so* $\mathtt{a} \Supset\!\Subset \mathtt{b}$*, in all implication space models.*

So we can see that, from the perspective of implication-space semantics, Elgin's notion of generalized identity coincides with sameness of content. It coincides with substitutability *salva consequentia* as premises and as conclusions in all models, and so merely in virtue of our semantic clauses for the logical vocabulary.

Our aim here is not to do justice to the debates in metaphysics that are related to the notions that we just recovered in implication-space semantics. Rather, we want to point out that these notions are isomorphic to notions that we can explain in terms of inclusion relations among implicational roles and, hence, ultimately in terms of substitution *salva consequentia* of bearers in implications, that is, in reason relations. In other words, we can introduce isomorphic notions by appealing just to substitutability *salva consequentia* in reason relations.

5.5.4 Nonlogical Role Inclusions

We have so far used implicational-role inclusions only to capture logical relations and concepts, namely familiar logics and (analogues of) concepts from analytic metaphysics. We want to end this chapter by considering an example of the possibility of codifying *non*logical concepts in implication-space semantics. In particular, we want to show how (a variant of) a recent inferentialist theory of predicates, due to Kai Tanter (2021), can be captured by implicational-role inclusions. For this subsection, we accordingly briefly lift the general restriction of our considerations to just the sentential level, and look at subsentential expressions. However, we will not consider quantifiers, and simply include schematic letters for singular terms to talk about open sentences.

Tanter (2021) offers an account of three very general kinds of conceptual connections, and he shows how to codify conceptual connections of these kinds in a sequent calculus. The three kinds of conceptual connections are, in our terminology:

1. Conceptual equivalence: F and G are such that, in virtue of the meanings of F and G, something is F if and only if it is G.
2. Species-genus relation: $F_1, \ldots, F_n$ and G are such that, in virtue of their meanings, being F_j is incompatible with being F_k for all $k \neq j$, and being any F_j is sufficient for being G.

3. Property-cluster relation: F and $G_1, ..., G_m$ are such that, in virtue of their meanings, being F is sufficient to be G_i for every i, and if something has all the properties $G_1, ..., G_m$, then it is F.

These three kinds of conceptual relations seem indeed pretty fundamental and powerful. These are the kinds of relation that allow us to build a Porphyrian tree, and do so with quite a bit of sophistication. We can, for example, codify the fact that, in virtue of the meanings of the terms, something is a computer program just in case it is a computer application. Using the second relation, we can codify that, in virtue of the meanings of the terms (let us assume), bacteria, archaea, and eukaryota are the most general species—namely the domains—of the genus of living beings. And the third kind of relations allows us to codify, in virtue of the meanings of the terms (let us assume), that something is a mammal just in case it is an animal and has hair. Inferential relations of these sorts are ubiquitous in ordinary language (colors, shapes, places...) and have accordingly loomed large in the history of thinking about the structure of concepts and reality, at least from Porphyry's *Isagoge* up through Linnaean taxonomy.

Tanter (2021) gives sequent rules for each of these three cases. We can formulate adjusted and slightly changed versions[25] of these rules, letting ν be schematic for singular terms and using double-lines to say that the sequent rule applies in both directions, so that one can not only infer the bottom sequent from the top sequents but also any top sequent from the bottom sequent. For conceptual equivalence the sequent rules are:

$$\frac{\Gamma, F\nu \succ \Delta}{\Gamma, G\nu \succ \Delta}\ \text{G/F-L} \qquad \frac{\Gamma \succ F\nu, \Delta}{\Gamma \succ G\nu, \Delta}\ \text{G/F-R}$$

This obviously ensures that F and G have the same premisory and conclusory implicational roles. So, we can formulate this in terms of implicational role inclusion as follows, using typewriter font for roles and superscripts for marking premisory and conclusory roles (with plus and minus, respectively):

G/F-L: $\mathtt{F}\nu^+ \preceq\succeq \mathtt{G}\nu^+$

G/F-R: $\mathtt{F}\nu^- \preceq\succeq \mathtt{G}\nu^-$

This says that, for all instances of these schematic sentences, the premisory roles of $F\nu$ and $G\nu$ include each other. And the same holds for their conclusory roles. Since mutual inclusion of implicational roles implies their identity, we could also express this as: $\mathtt{F}\nu^+ = \mathtt{G}\nu^+$ and $\mathtt{F}\nu^- = \mathtt{G}\nu^-$. If we

used single-line sequent rules, then the inclusion relations might hold in only one direction.

The rules for the Species-Genus relation are as follows, where F_i and F_j stand for arbitrary but distinct Fs such that the sequence $F_1, ..., F_i$ is like $F_1, ..., F_n$ but without F_j.

$$\frac{\Gamma \succ F_1\nu, ..., F_n\nu, \Delta}{\Gamma \succ G\nu, \Delta}\ \text{G/FR} \qquad \frac{\Gamma, F_1\nu \succ \Delta \quad \cdots \quad \Gamma, F_n\nu \succ \Delta}{\Gamma, G\nu \succ \Delta}\ \text{G/FL}$$

$$\frac{\Gamma \succ G\nu, \Delta \qquad \Gamma, F_1\nu \succ \Delta \quad \cdots \quad \Gamma, F_i\nu \succ \Delta}{\Gamma \succ F_j\nu, \Delta}\ \text{RFj} \qquad \frac{\Gamma, G\nu \succ F_1\nu, ..., F_i\nu, \Delta}{\Gamma, F_j\nu \succ \Delta}\ \text{LFj}$$

G/FR says that it follows that something falls under one or another of the species just in case it follows that it falls under the genus. G/FL says that something follows from an object falling under any species of a genus just in case it follows from the object falling under that genus. Note that G/FR and G/FL are what you would expect if $G\nu$ is the disjunction of all the $F_k\nu$. We can express these rules in implicational role inclusion terms as follows:

G/FR: $\mathtt{F_1}\nu^- \sqcup ... \sqcup \mathtt{F_n}\nu^- \preceq\succeq \mathtt{G}\nu^-$

G/FL: $\mathtt{F_1}\nu^+ \sqcap ... \sqcap \mathtt{F_n}\nu^+ \preceq\succeq \mathtt{G}\nu^+$

The rule LFj says that all the species are mutually exclusive: It follows from something falling under the genus, given premises Γ, that either it falls under one of the species $F_1, ..., F_i$ or one of the sentences Δ holds if and only if it follows from it falling under the remaining species F_j, given premises Γ, that one of the sentences Δ holds. Notice that RFj and LFj are what you would expect if F_j is equivalent to $G\nu \wedge \neg F_1\nu \wedge ... \wedge \neg F_i\nu$. We can formulate these rules in role inclusion terms as follows:

RFj: $\mathtt{G}\nu^- \sqcap \mathtt{F_1}\nu^+ \sqcap ... \sqcap \mathtt{F_i}\nu^+ \preceq\succeq \mathtt{F_j}\nu^-$

LFj: $\mathtt{G}\nu^+ \sqcup \mathtt{F_1}\nu^- \sqcup ... \sqcup \mathtt{F_i}\nu^- \preceq\succeq \mathtt{F_j}\nu^+$

The rules for the Property-Cluster relation are the following:

$$\frac{\Gamma \succ G_1\nu, \Delta \quad \cdots \quad \Gamma \succ G_m\nu, \Delta}{\Gamma \succ F\nu, \Delta}\ \text{F/GR} \qquad \frac{\Gamma, G_1\nu, ..., G_m\nu \succ \Delta}{\Gamma, F\nu \succ \Delta}\ \text{F/GL}$$

These are the rules one would expect if one wanted to say that $F\nu$ has the same inferential role as $G_1\nu \wedge ... \wedge G_m\nu$. For F/GL corresponds to the

usual conjunction left-rule and its inverse, and F/GR corresponds to the usual conjunction right-rule and its inverse. We can formulate these rules in role inclusion terms as follows:

F/GR: $\mathtt{G}_1\nu^- \sqcap \ldots \sqcap \mathtt{G}_\mathtt{m}\nu^- \preceq\succeq \mathtt{F}\nu^-$

F/GL: $\mathtt{G}_1\nu^+ \sqcup \ldots \sqcup \mathtt{G}_\mathtt{m}\nu^+ \preceq\succeq \mathtt{F}\nu^+$

One can read the role inclusions as holding *de jure*, and then they say, for example, that if you want to have concepts of genus and species, then you need to define these concepts as given by the appropriate role inclusions. But we can also read the role inclusions, as it were, *a posteriori* or *de facto*, and then they say, for example, that if you find concepts whose roles are related in the given ways, then you have found yourself a genus and some species concepts. The formalism is, of course, neutral on which of these readings is more appropriate (if either).

Like the sequent rules and the implication-role inclusions for logical connectives, we can interpret the sequent rules for these conceptual connections in truth-maker theory, and doing so might add a potentially illuminating perspective. Interpreted in truth-maker theory, the rule G/F-R says that all instantiations for ν of $F\nu$ and $G\nu$ have the same falsity-makers. G/F-L says that, for every instantiation of ν, $F\nu$ and $G\nu$ have the same truth-makers. So given these rules, $F\nu$ and $G\nu$ always represent the same worldly proposition.

Turning to the rules for species-genus relations, the rule G/FR says that, for every instantiation of ν, a falsity-maker of $G\nu$ is a fusion of falsity-makers of each of $F_1\nu, \ldots, F_n\nu$. We could express this by saying that to fail to fall under a genus is to fail to fall under any of its species. The rule G/FL says that, for every instantiation of ν, the truth-makers of $G\nu$ are the union of truth-makers for any of $F_1\nu, \ldots, F_n\nu$. We could express this by saying that to fall under a genus is to fall under one of its species. The RFj rule says that, for every instantiation of ν, a falsity-maker of $F_j\nu$ is either a falsity maker of $G\nu$ or a truth-maker of one of the $F_1\nu, \ldots, F_i\nu$. We could express this by saying that to fail to fall under a particular species is either to fail to fall under the genus or to belong to one of the other species under that genus. The LFj rule says that, for every instantiation of ν, a truth-maker of $F_j\nu$ is a fusion of a truth-maker of $G\nu$ and falsity-makers for each of the $F_1\nu, \ldots, F_i\nu$. We could express that by saying that to fall under a particular species is to fall under the genus and fail to fall under any of its other species.

Finally, regarding the rules for property-cluster relations, the rule F/GR says that, for every instantiation of ν, a falsity-maker of $F\nu$ is anything that is a falsity-maker for one of $G_1\nu, \ldots, G_m\nu$. We could express this by saying that to fail to fall under a property-cluster concept is for something to fail to

have one (or more) of the properties in the cluster that defined the property-cluster concept. And the rule F/GL says that, for every instantiation of ν, a truth-maker of $F\nu$ is a fusion of truth-makers for each of the $G_1\nu, \ldots, G_m\nu$. We could express that by saying that to fall under a property-cluster concept is to have all the properties in the cluster.

To sum up, the metalanguage of implicational-role inclusions allows us to formulate what reason relations must be like for some concepts in these reason relations to stand in the relations of conceptual equivalence, the species-genus relation, or the property-cluster relation. The correspondence between implication-space semantics, sequent calculi, and truth-maker theory applies to these cases in the same way in which it applies to logical vocabulary. This illustrates how implication-space semantics can codify the kind of chain of being familiar since Porphyry's *Isagoge*, at least until Linnaean taxonomy, in an abstract metalanguage of reason relations that captures the structure that is shared between more concrete characterizations of reason relations in (the right kind of) pragmatic-normative and semantic-representationalist metavocabularies.

5.6 Conclusion

In this chapter, we explicated reason relations in a metavocabulary of implicational roles. This explication of reason relations is intrinsic because it does not appeal to anything that is not already given in the reason relations themselves. In particular, it does not appeal to norms governing discursive acts or modal relations among worldly states. Rather, implication-space semantics takes implication relations among some particular items and abstracts from them the implicational role of these items. The pragmatic-normative account of reason relations in terms of the sequent calculus NMMS from Chapter Three and the semantic-representationalist account in terms of truth-makers from Chapter Four can both be interpreted in implication-space semantics. Indeed, we can use the same implicational roles to interpret the role of a sentence in the pragmatic-normative account and to interpret the worldly proposition that the sentence represents in the semantic-representationalist account. We have thus arrived at an abstract theory of the rational forms that are shared by sentences and the worldly propositions that they represent: a theory that abstracts away from the different ways in which the rational form can be realized in the matter of discursive acts or the matter of worldly states. We have arrived at a pure theory of rational forms.

Let us rehearse the main points in the development of implication-space semantics. We start with a space of bearers of implicational roles and an implication relation between sets of these bearers. We call this an implication frame. We then consider under which additions a candidate

implication would turn into a good implication, or remain good (if it already is a good implication). We call the set of these additions the "range of subjunctive robustness" of the candidate implication. The implicational role of an implication is then the class of (sets of) implications with the same range of subjunctive robustness. The premisory role of a bearer is the role of the implication in which the bearer occurs only as a premise. The conclusory role of a bearer is the role of the implication in which the bearer occurs only as a conclusion. And the implicational role of a bearer is the pair of its premisory and its conclusory roles.

Any pair that could be the implicational role of a bearer is a conceptual content. And we can interpret languages by assigning to each sentence of the language such a conceptual content. A model is a set of contents that interpret the sentences of a given language. And a model-theoretic implication holds, in a model, just in case all implications in the role of the implication are good implications (in the implication frame of the contents of the model). Finally, a model-theoretic implication holds in a set of models if it holds in each of the models.

The mathematical structure that is necessary and sufficient to define an implication frame is a commutative monoid, in which the monoid set is a set of pairs, together with a bipartition of the monoid set. The elements of the monoid set are candidate implications, and the operation on them allows us to combine implications with each other. Ranges of subjunctive robustness correspond to the set of candidate implications that one can combine with a target implication so that the result is in the particular subset of the monoid set. Hence, the fundamental structure of reason relations is the structure of a commutative monoid in which the monoid set is a set of pairs, together with a partition of the monoid set.

We define the implicational roles of logically complex sentences in terms of the implicational roles of their constituent sentences. And we have shown that if we define the connectives in a particular way, and we look only at implication-space models in which the material implications of some base $\mathfrak{B}$ hold, then $\mathsf{NMMS}_{\mathfrak{B}}$ is sound and complete with respect to the consequence relation of implication-space semantics over the set of these models, even if the reason relations of the base are radically open-structured. If we define the connectives in other ways and allow for failures of Contraction, we can capture $\mathsf{NMMS}_{\mathfrak{B}}^{/ctr}$ and multiplicative additive linear logic in implication-space semantics.

Finally, we can define inclusion relations among implicational roles by saying that one role includes another if the range of subjunctive robustness of the first role is a subset of the range of subjunctive robustness of the second role. This notion of implicational role inclusion allows us to recover within the framework of implication-space semantics all of: the logic of

paradox, strong Kleene logic, strict-tolerant logic, tolerant-strict logic, and several relations and notions from analytic metaphysics.

Thus, once we reach the level of pure rational forms, all the familiar logics and notions just mentioned fall into place as particular aspects or variations on relations among these rational forms—that is, relations among implicational roles. For instance, the ideas of truth-value gluts and gaps and the ideas of factual equivalence, essence, and grounding can all be understood in terms of implicational role inclusions. And the connectives of linear logic can also be defined in terms of implicational roles.

We have shown how one can use implication-space semantics to formulate a broad variety of familiar logics and formal theories more generally. We hypothesize, however, that there are many more such theories, and that one can extend and generalize our results in many directions. One promising project is to extend implication-space semantics to cover more logical vocabulary, such as quantifiers, standard modal operators, and other kinds of conditionals. Another is to investigate where and how the connections between implication-space semantics and the sequent calculus and truth-maker theory may break down. Yet another is to recover further familiar logics in the framework of implicational-role inclusions, with Angell's logic of analytic contaiment, the logic of first-degree entailment, and connexive logics suggesting themselves as proximal targets. These and many similar projects seem to us to be tractable and potentially fruitful research projects, which we take to support the idea that implication-space semantics can provide the tools for moving many debates forward.

Implication-space semantics is a formal inferentialist semantics. For according to implication-space semantics, the contents of sentences are the roles they play in implications, that is, the roles they play in reason relations. The semantic clauses for the logical connectives show in what sense such an inferentialist semantic theory is compositional, namely in the sense that the implicational roles of logically complex sentences are defined in terms of the implicational roles of their constituents. The semantic theory is, nevertheless, holistic because the role of an implication is defined in terms of which additions of further premises and conclusions would yield a good implication. In this way, the implicational role of an implication points beyond the implication's constituents and to other implications. Hence, the implication-space semantics of the logical connectives gives us a concrete example of a semantics that is holistic but also compositional. Moreover, implication-space semantics shows how an inferentialist semantics can capture open reason relations in a formally rigorous way, thus codifying substructural consequence relations. Finally, the many connections between implication-space semantics and extant logical and metaphysical theories suggest that implication-space semantics

may serve as a unifying framework for many different approaches, from normative bilateralism, to truth-maker theory, three-valued logics, linear logic, the logics of grounding and essence and generalized identity. We take this expressive flexibility and power to be indirect evidence that implication-space semantics captures the fundamental structure of reason relations and, hence, the fundamental structure of conceptual content.

5.7 Appendix

5.7.1 Roles and Ranges of Subjunctive Robustness

We start by listing some simple propositions about roles and ranges of subjunctive robustness, which are useful when one is trying to find one's way into implication-space semantics.

Proposition 102. *If* $\alpha = \mathtt{RSR}(\mathtt{RSR}(\alpha))$ *and* $\beta = \mathtt{RSR}(\mathtt{RSR}(\beta))$ *and* α *and* β *both are members of some role* $\mathcal{R}$*, then* $\alpha = \beta$*.*

Proof. Clearly, if $\alpha = \mathtt{RSR}(\mathtt{RSR}(\alpha))$ and $\beta = \mathtt{RSR}(\mathtt{RSR}(\beta))$ and $\mathtt{RSR}(\alpha) = \mathtt{RSR}(\beta)$, then $\alpha = \beta$. But, by the definition of roles, if α and β both are members of some role $\mathcal{R}$, then $\mathtt{RSR}(\alpha) = \mathtt{RSR}(\beta)$. ■

Proposition 103. $\mathtt{RSR}(\mathcal{R}(\alpha)) = \mathtt{RSR}(\alpha)$.

Proof. By the definition of implicational roles, $\forall x \in \mathcal{R}(\alpha)$ $(\mathtt{RSR}(\alpha) = \mathtt{RSR}(x))$. But $\mathtt{RSR}(X) = \bigcap_{x \in X} \mathtt{RSR}(x)$ and, hence, $\mathtt{RSR}(\mathcal{R}(\alpha)) = \mathtt{RSR}(x)$, for any $x \in \mathcal{R}(\alpha)$. So $\mathtt{RSR}(\alpha) = \mathtt{RSR}(\mathcal{R}(\alpha))$. ■

Proposition 104. $\mathcal{R}(\mathcal{R}(\alpha)) = \mathcal{R}(\alpha)$.

Proof. By the definition of roles, $\mathcal{R}(\mathcal{R}(\alpha))$ is the set $\{x \mid \mathtt{RSR}(\mathcal{R}(\alpha)) = \mathtt{RSR}(x)\}$. By Proposition 103, this is the set $\{x \mid \mathtt{RSR}(\alpha) = \mathtt{RSR}(x)\}$. But that is $\mathcal{R}(\alpha)$. ■

5.7.2 Implication-Space Semantics and NMMS

We now move on to proving that the results of taking an adjunction or symjunction of two implicational roles is independent of the way in which these roles are represented. That is, if two items have the same implicational roles, then we do not get different results for the two items when we compute adjunctions and symjunctions of their roles.

Proposition 105. *If* $\mathcal{R}(F) = \mathcal{R}(F')$*, then* $\mathcal{R}(F) \sqcup \mathcal{R}(G) = \mathcal{R}(F') \sqcup \mathcal{R}(G)$*.*

Proof. Suppose that $\mathcal{R}(F) = \mathcal{R}(F')$. To show that $\mathcal{R}(F) \sqcup \mathcal{R}(G) = \mathcal{R}(F') \sqcup \mathcal{R}(G)$ it suffices to show that $\mathtt{RSR}(\{\langle f^+ \cup g^+, f^- \cup g^- \rangle \mid \langle f^+, f^- \rangle \in F, \langle g^+, g^- \rangle \in G\})$ is the same as $\mathtt{RSR}(\{\langle f'^+ \cup g^+, f'^- \cup g^- \rangle \mid \langle f'^+, f'^- \rangle \in F', \langle g^+, g^- \rangle \in G\})$. So, suppose that $\langle x, y \rangle \in \mathtt{RSR}(\mathcal{R}(F) \sqcup \mathcal{R}(G))$. Then, for every $\langle g^+, g^- \rangle \in G$, we have $\langle x \cup g^+, y \cup g^- \rangle \in \mathtt{RSR}(F)$. Since $\mathtt{RSR}(F) = \mathtt{RSR}(F')$, for every $\langle g^+, g^- \rangle \in G$, we have $\langle x \cup g^+, y \cup g^- \rangle \in \mathtt{RSR}(F')$. So $\langle x, y \rangle \in \mathtt{RSR}(\mathcal{R}(F') \sqcup \mathcal{R}(G))$. And the same reasoning works in the other direction. Therefore, $\langle x, y \rangle \in \mathtt{RSR}(\mathcal{R}(F) \sqcup \mathcal{R}(G))$ just in case $\langle x, y \rangle \in \mathtt{RSR}(\mathcal{R}(F') \sqcup \mathcal{R}(G))$. So $\mathcal{R}(F) \sqcup \mathcal{R}(G) = \mathcal{R}(F') \sqcup \mathcal{R}(G)$. ■

Proposition 106. *If* $\mathcal{R}(F) = \mathcal{R}(F')$, *then* $\mathcal{R}(F) \sqcap \mathcal{R}(G) = \mathcal{R}(F') \sqcap \mathcal{R}(G)$.

Proof. If $\mathcal{R}(F) = \mathcal{R}(F')$, then $F' \approx F$. Now, $\mathcal{R}(F) \sqcap \mathcal{R}(G) = \mathcal{R}(F \cup G)$, and the latter is the set of sets of implications such that each of the sets has the same range of subjunctive robustness as $F \cup G$. So it suffices to show that, for all $\langle X, Y \rangle$, we have $\forall \langle \Gamma, \Delta \rangle \in F \cup G\ (\langle \Gamma \cup X, \Delta \cup Y \rangle \in \mathbb{I})$ iff $\forall \langle \Gamma', \Delta' \rangle \in F' \cup G\ (\langle \Gamma' \cup X, \Delta' \cup Y \rangle \in \mathbb{I})$. This holds because $F' \approx F$ entails that $\forall \langle \Gamma, \Delta \rangle \in F\ (\langle \Gamma \cup X, \Delta \cup Y \rangle \in \mathbb{I})$ iff $\forall \langle \Gamma', \Delta' \rangle \in F'\ (\langle \Gamma' \cup X, \Delta' \cup Y \rangle \in \mathbb{I})$. ■

The following lemma about the connection between adjunction and symjunction will prove useful below.

Lemma 107. $\bigcup((\mathtt{a} \sqcap \mathtt{b}) \sqcup \mathtt{c}) \subseteq \mathbb{I}$ *just in case* $\bigcup(\mathtt{a} \sqcup \mathtt{c}) \subseteq \mathbb{I}$ *and* $\bigcup(\mathtt{b} \sqcup \mathtt{c}) \subseteq \mathbb{I}$.

Proof. It follows from the definition of symjunction that $\mathtt{RSR}(\mathtt{a} \sqcap \mathtt{b})$ is the intersection of $\mathtt{RSR}(\mathtt{a})$ and $\mathtt{RSR}(\mathtt{b})$. Hence, for all $c \in \mathtt{c}$, $c \subseteq \mathtt{RSR}(\mathtt{a} \sqcap \mathtt{b})$ if and only if $c \subseteq \mathtt{RSR}(\mathtt{a})$ and $c \subseteq \mathtt{RSR}(\mathtt{b})$. But, in general, for any role $\mathtt{z}$, we have $\bigcup(\mathtt{z} \sqcup \mathtt{c}) \subseteq \mathbb{I}$ just in case, for all $c \in \mathtt{c}$, $c \subseteq \mathtt{RSR}(\mathtt{z})$. So, $\bigcup((\mathtt{a} \sqcap \mathtt{b}) \sqcup \mathtt{c}) \subseteq \mathbb{I}$ just in case, for all $c \in \mathtt{c}$, $c \subseteq \mathtt{RSR}(\mathtt{a} \sqcap \mathtt{b})$. And the latter holds just in case, for all $c \in \mathtt{c}$, $c \subseteq \mathtt{RSR}(\mathtt{a})$ and $c \subseteq \mathtt{RSR}(\mathtt{b})$. The latter holds if and only if $\bigcup(\mathtt{a} \sqcup \mathtt{c}) \subseteq \mathbb{I}$ and $\bigcup(\mathtt{b} \sqcup \mathtt{c}) \subseteq \mathbb{I}$. ■

We now turn to the relation between implication-space semantics and the sequent calculus NMMS from Chapter Three.

Proposition 108. *If there is an* NMMS *rule application with top sequents* $\Gamma_1 \succ \Delta_1$, ..., $\Gamma_n \succ \Delta_n$ *and bottom sequent* $\Gamma_0 \succ \Delta_0$, *then, for all implication–space models,* $\Gamma_0 \overset{\mathcal{M}}{\mid\!\sim} \Delta_0$ *holds just in case all of* $\Gamma_1 \overset{\mathcal{M}}{\mid\!\sim} \Delta_1$ *and ... and* $\Gamma_n \overset{\mathcal{M}}{\mid\!\sim} \Delta_n$ *hold.*

Proof. We do only the cases for negation and the conditional, as conjunction and disjunction can be defined in terms of negation and the conditional (and also the proofs work analogously for them). We take each of the rules in turn, always letting $[\![A]\!] = \langle \mathtt{a}^+, \mathtt{a}^- \rangle$ and $[\![B]\!] = \langle \mathtt{b}^+, \mathtt{b}^- \rangle$ and $[\![\Gamma]\!] = \mathtt{P} = \{\mathtt{p}_0, ..., \mathtt{p}_n\}$ and $[\![\Delta]\!] = \mathtt{C} = \{\mathtt{c}_0, ..., \mathtt{c}_m\}$.

[L¬] concludes $\Gamma, \neg A \succ \Delta$ from $\Gamma \succ A, \Delta$. By our semantic clauses for negation, $[\![\neg A]\!] = \langle (\neg \mathtt{a})^+, (\neg \mathtt{a})^- \rangle = \langle \mathtt{a}^-, \mathtt{a}^+ \rangle$. Suppose that $\Gamma \overset{\mathcal{M}}{\mid\!\sim} A, \Delta$ and, hence, $\mathtt{P} \Vdash\!\sim \mathtt{a}, \mathtt{C}$. That is, $\bigcup(\bigsqcup_{i=0}^{n} \mathtt{p}_i^+ \sqcup \bigsqcup_{j=0}^{m} \mathtt{c}_j^- \sqcup \mathtt{a}^-) \subseteq \mathbb{I}$. But $\mathtt{a}^- = (\neg \mathtt{a})^+$. So, $\bigcup(\bigsqcup_{i=0}^{n} \mathtt{p}_i^+ \sqcup \bigsqcup_{j=0}^{m} \mathtt{c}_j^- \sqcup (\neg \mathtt{a})^+) \subseteq \mathbb{I}$. Hence, $\mathtt{P}, \neg \mathtt{a} \Vdash\!\sim \mathtt{C}$ and, so, $\Gamma, \neg A \overset{\mathcal{M}}{\mid\!\sim} \Delta$. And the same reasoning works in the reverse direction as well.

[R¬] concludes $\Gamma \succ \neg A, \Delta$ from $\Gamma, A \succ \Delta$. Suppose that $\Gamma, A \overset{\mathcal{M}}{\mid\!\sim} \Delta$ and, hence, $\mathtt{P}, \mathtt{a} \Vdash\!\sim \mathtt{C}$. That is, $\bigcup(\bigsqcup_{i=0}^{n} \mathtt{p}_i^+ \sqcup \mathtt{a}^+ \sqcup \bigsqcup_{j=0}^{m} \mathtt{c}_j^-) \subseteq \mathbb{I}$. But $\mathtt{a}^+ = (\neg \mathtt{a})^-$. So, $\bigcup(\bigsqcup_{i=0}^{n} \mathtt{p}_i^+ \sqcup (\neg \mathtt{a})^- \sqcup \bigsqcup_{j=0}^{m} \mathtt{c}_j^-) \subseteq \mathbb{I}$. Hence, $\mathtt{P} \Vdash\!\sim \neg \mathtt{a}, \mathtt{C}$ and, so, $\Gamma \overset{\mathcal{M}}{\mid\!\sim} \neg A, \Delta$. And the same reasoning works in the reverse direction as well.

[L→] concludes $\Gamma, A \to B \succ \Delta$ from $\Gamma \succ \Delta, A$ and $\Gamma, B \succ \Delta$ and $\Gamma, B \succ \Delta, A$. By our semantic clauses for the conditional $[\![A \to B]\!] = \langle (\mathtt{a} \to \mathtt{b})^+, (\mathtt{a} \to \mathtt{b})^- \rangle = \langle \mathtt{a}^- \sqcap \mathtt{b}^+ \sqcap (\mathtt{a}^- \sqcup \mathtt{b}^+), \mathtt{a}^+ \sqcup \mathtt{b}^- \rangle$. Suppose that $\Gamma, A \to B \overset{\mathcal{M}}{\mid\!\sim} \Delta$ and, hence, $\mathtt{P}, \mathtt{a} \to \mathtt{b} \Vdash\!\sim \mathtt{C}$. That is, $\bigcup(\bigsqcup_{i=0}^{n} \mathtt{p}_i^+ \sqcup (\mathtt{a}^- \sqcap \mathtt{b}^+ \sqcap (\mathtt{a}^- \sqcup \mathtt{b}^+)) \sqcup \bigsqcup_{j=0}^{m} \mathtt{c}_j^-) \subseteq \mathbb{I}$. By Lemma 107, that holds just in case $\bigcup(\bigsqcup_{i=0}^{n} \mathtt{p}_i^+ \sqcup \mathtt{a}^- \sqcup \bigsqcup_{j=0}^{m} \mathtt{c}_j^-) \subseteq \mathbb{I}$ and $\bigcup(\bigsqcup_{i=0}^{n} \mathtt{p}_i^+ \sqcup \mathtt{b}^+ \sqcup \bigsqcup_{j=0}^{m} \mathtt{c}_j^-) \subseteq \mathbb{I}$ and $\bigcup(\bigsqcup_{i=0}^{n} \mathtt{p}_i^+ \sqcup \mathtt{a}^- \sqcup \mathtt{b}^+ \sqcup \bigsqcup_{j=0}^{m} \mathtt{c}_j^-) \subseteq \mathbb{I}$. Which holds just in case $\mathtt{P} \Vdash\!\sim \mathtt{a}, \mathtt{C}$ and $\mathtt{P}, \mathtt{b} \Vdash\!\sim \mathtt{C}$ and $\mathtt{P}, \mathtt{b} \Vdash\!\sim \mathtt{a}, \mathtt{C}$. So, $\Gamma \overset{\mathcal{M}}{\mid\!\sim} A, \Delta$ and $\Gamma, B \overset{\mathcal{M}}{\mid\!\sim} \Delta$ and $\Gamma, B \overset{\mathcal{M}}{\mid\!\sim} A, \Delta$. And the same reasoning works in the reverse direction as well.

[R→] concludes $\Gamma \succ A \to B, \Delta$ from $\Gamma, A \succ B, \Delta$. Suppose that $\Gamma, A \overset{\mathcal{M}}{\mid\!\sim} B, \Delta$ and, hence, $\mathtt{P}, \mathtt{a} \Vdash\!\sim \mathtt{b}, \mathtt{C}$. That is, $\bigcup(\bigsqcup_{i=0}^{n} \mathtt{p}_i^+ \sqcup \mathtt{a}^+ \sqcup \mathtt{b}^- \sqcup \bigsqcup_{j=0}^{m} \mathtt{c}_j^-) \subseteq \mathbb{I}$. But $\mathtt{a}^+ \sqcup \mathtt{b}^- = (\mathtt{a} \to \mathtt{b})^-$. So, $\bigcup(\bigsqcup_{i=0}^{n} \mathtt{p}_i^+ \sqcup (\mathtt{a} \to \mathtt{b})^- \sqcup \bigsqcup_{j=0}^{m} \mathtt{c}_j^-) \subseteq \mathbb{I}$. Hence, $\mathtt{P} \Vdash\!\sim \mathtt{a} \to \mathtt{b}, \mathtt{C}$ and, so, $\Gamma \overset{\mathcal{M}}{\mid\!\sim} A \to B, \Delta$. And the same reasoning works in the reverse direction as well. ■

As a matter of terminology we let b be the set of models that are fit for the base $\mathfrak{B}$ (see the main text for the definition of fitness of a model for a base).

Theorem 109. *For any base* $\mathfrak{B} = \left\langle \mathfrak{L}_{\mathfrak{B}}, \mathrel{|\!\sim}_{\mathfrak{B}} \right\rangle$ *and sentences* $\Gamma, \Delta \subseteq \mathfrak{L}$ *in the logically extended lexicon of the base,* $\Gamma \overset{b}{\mathrel{|\!\sim}} \Delta$ *if and only if* $\Gamma \succ \Delta$ *is derivable in* $\mathsf{NMMS}_{\mathfrak{B}}$.

Proof. By Proposition 108, implications in implication-space models are closed under the rules of NMMS. And if $\Gamma \mathrel{|\!\sim}_{\mathfrak{B}} \Delta$ in the base consequence relation, then, by definition, $\Gamma \mathrel{|\!\sim} \Delta$ holds in all models fit for $\mathfrak{B}$. So if $\Gamma \succ \Delta$ is derivable in $\mathsf{NMMS}_{\mathfrak{B}}$, then $\Gamma \overset{b}{\mathrel{|\!\sim}} \Delta$. For the other direction, if $\Gamma \succ \Delta$ is not derivable in $\mathsf{NMMS}_{\mathfrak{B}}$, then a proof-search on $\Gamma \succ \Delta$ yields atomic sequents that are not in the base consequence relation $\mathrel{|\!\sim}_{\mathfrak{B}}$. We can construct an implication-space b-model in which these sequents fail. In fact, the model with the implication frame $\left\langle \mathfrak{L}_{\mathfrak{B}}, \mathrel{|\!\sim}_{\mathfrak{B}} \right\rangle$ and the interpretation $[\![\phi]\!]^{\mathcal{M}} = \mathcal{R}(\phi)$ for all $\phi \in \mathfrak{L}$ can serve as a counterexample. ■

5.7.3 Implication-Space Semantics and Truth-Maker Theory

We now turn to the relation between implication-space semantics and the truth-maker theory from Chapter Four.

Theorem 110. *Let* $\mathcal{M}$ *be the implication-space model defined by the implication frame,* $\langle \mathtt{B}, \mathbb{I}_{\overline{\diamond}} \rangle$*, of the modalized state space* $\langle S, S^{\diamond}, \sqsubseteq \rangle$ *and the interpretation function such that* $[\![x]\!]^{\mathcal{M}} = \mathcal{R}(x)$ *for all* $x \in \mathtt{B}$*. Then* $P \mathrel{|\!\!\models}_{PI} C$ *holds in the modalized state space if and only if* $P \overset{\mathcal{M}}{\mathrel{|\!\sim}} C$.

Proof. Let $P = \{p_0, ..., p_n\}$ and $C = \{c_0, ..., c_m\}$, and suppose that $P \mathrel{|\!\!\models}_{PI} C$ holds in a modalized state space, i.e., any fusion of verifiers for each $\langle p_i^+, p_i^- \rangle \in P$ and falsifiers for each $\left\langle \mathtt{c}_j^+, c_j^- \right\rangle \in C$ is an impossible state. So, by the definition of the implication frame of a modalized state space, $\langle P, C \rangle \in \mathbb{I}_{\overline{\diamond}}$. By the reasoning from the proof of Proposition 74 above, $\langle P, C \rangle \in \mathbb{I}_{\overline{\diamond}}$ if and only if $P \overset{\mathcal{M}}{\mathrel{|\!\sim}} C$. ■

Proposition 111. *If the implication-space model* $\mathcal{M}$ *and truth-maker model* $\mathcal{M}'$ *are parallel, then* $\Gamma \overset{\mathcal{M}}{\mathrel{|\!\sim}} \Delta$ *just in case* $\Gamma \mathrel{|\!\!\models}_{TM} \Delta$ *in* $\mathcal{M}'$.

Proof. Let $\mathcal{M}''$ be the model with the same implication frame, $\langle \mathtt{B}, \mathbb{I}_{\overline{\diamond}} \rangle$, as $\mathcal{M}$ but an interpretation function that interprets worldly propositions by themselves, namely the interpretation such that $[\![x]\!]^{\mathcal{M}''} = \mathcal{R}(x)$ for all $x \in \mathtt{B}$. And let f be a function that takes sets of sentences of $\mathfrak{L}$ to the sets of their interpretations in $\mathcal{M}'$. We now show that implications in $\mathcal{M}$ and

$\mathcal{M}'$ both coincide with implications in $\mathcal{M}''$, thus using $\mathcal{M}''$ as our "middle man."

$\mathcal{M}$ and $\mathcal{M}''$: If $[\![\Gamma]\!]^{\mathcal{M}} = [\![f(\Gamma)]\!]^{\mathcal{M}''}$ and $[\![\Delta]\!]^{\mathcal{M}} = [\![f(\Delta)]\!]^{\mathcal{M}''}$, then $\Gamma \overset{\mathcal{M}}{\mid\!\sim} \Delta$ just in case $f(\Gamma) \overset{\mathcal{M}''}{\mid\!\sim} f(\Delta)$. For, $\mathcal{M}$ and $\mathcal{M}''$ differ only in the language and interpretation function. Since the interpretations of $\mathcal{M}$ and $\mathcal{M}'$ are parallel, $[\![A]\!]^{\mathcal{M}} = |A|^{\mathcal{M}'} = \mathcal{R}(A)$ and so $\mathcal{R}(|A|^{\mathcal{M}'}) = \mathcal{R}(\mathcal{R}(A))$. Moreover, we defined $\mathcal{M}''$ such that $\left[\!\!\left[|A|^{\mathcal{M}'}\right]\!\!\right]^{\mathcal{M}''} = \mathcal{R}(|A|^{\mathcal{M}'})$. It is easy to see, however, that $\mathcal{R}(A) = \mathcal{R}(\mathcal{R}(A))$ (Proposition 104). Hence, $[\![A]\!]^{\mathcal{M}} = \mathcal{R}(|A|^{\mathcal{M}'}) = [\![f(\{A\})]\!]^{\mathcal{M}''}$, which implies that $[\![\Gamma]\!]^{\mathcal{M}} = [\![f(\Gamma)]\!]^{\mathcal{M}''}$ and $[\![\Delta]\!]^{\mathcal{M}} = [\![f(\Delta)]\!]^{\mathcal{M}''}$. Therefore, $\Gamma \overset{\mathcal{M}}{\mid\!\sim} \Delta$ just in case $f(\Gamma) \overset{\mathcal{M}''}{\mid\!\sim} f(\Delta)$.

$\mathcal{M}'$ and $\mathcal{M}''$: To show that $\Gamma \underset{TM}{\models\!\!=} \Delta$ in $\mathcal{M}'$ just in case $f(\Gamma) \overset{\mathcal{M}''}{\mid\!\sim} f(\Delta)$, we first note that, by Theorem 79, $P \underset{PI}{\Vdash\!\!=} C$ holds in the modalized state space of all three models if and only if $P \overset{\mathcal{M}''}{\mid\!\sim} C$. Hence, it suffices to note that $[\![\Gamma]\!]^{\mathcal{M}'} = f(\Gamma)$ and $[\![\Delta]\!]^{\mathcal{M}'} = f(\Delta)$ holds by the definition of f. ■

Proposition 112. *If there is an implication-space model $\mathcal{M}$ such that $\overset{\mathcal{M}}{\mid\!\sim} = X$, then there is a truth-maker model $\mathcal{M}'$ such that $\underset{TM}{\models\!\!=} X$.*

Proof. Suppose that there is an implication-space model $\mathcal{M}$ such that $\overset{\mathcal{M}}{\mid\!\sim} = X$. The modalized state space of the desired model $\mathcal{M}'$ is the space $\langle S, S^{\diamond}, \sqsubseteq\rangle$ defined as follows: Let S be the set of all pairs such that $\langle \Gamma^{At}, \Delta^{At}\rangle$ with Γ^{At} and Δ^{At} being sets of atomic sentences of the language, $\mathfrak{L}$, that $\mathcal{M}$ interprets. Let $\langle \Gamma^{At}, \Delta^{At}\rangle \sqsubseteq \langle \Theta^{At}, \Lambda^{At}\rangle$ if and only if $\Gamma^{At} \subseteq \Theta^{At}$ and $\Theta^{At} \subseteq \Lambda^{At}$. Let $S^{\diamond}$ be the set such that $\langle \Gamma^{At}, \Delta^{At}\rangle \in S^{\diamond}$ if and only if $\langle \Gamma^{At}, \Delta^{At}\rangle \notin X$. The interpretation function of $\mathcal{M}'$ is the following: For all atomic sentences ϕ in $\mathfrak{L}$, let $|\phi|^{+} = \{\langle\{\phi\}, \emptyset\rangle\}$ and $|\phi|^{-} = \{\langle\emptyset, \{\phi\}\rangle\}$. We extend this interpretation of atomic sentences by the usual truth-maker clauses for the logical vocabulary.

Notice that since fusion is the least upper bound with respect to parthood, the fusion of $\langle \Gamma^{At}, \Delta^{At}\rangle$ and $\langle \Theta^{At}, \Lambda^{At}\rangle$ is $\langle \Gamma^{At} \cup \Theta^{At}, \Delta^{At} \cup \Lambda^{At}\rangle$. Hence, $\langle \Gamma^{At}, \Delta^{At}\rangle \notin S^{\diamond}$ just in case the fusion of the truth-makers for each element in Γ^{At} and the falsity-makers for each element in Δ^{At} is an impossible state. Hence, for atomic sentences, $\langle \Gamma^{At}, \Delta^{At}\rangle \notin S^{\diamond}$ just in case $\Gamma^{At} \underset{TM}{\models\!\!=} \Delta^{At}$.

We now consider the semantic clauses. We do just the case for negation and conjunction, as the other connectives can be defined in terms of these. We argue by induction on the number of connectives in a candidate implication.

Negation: Suppose that $\Gamma \overset{\mathcal{M}}{\mid\!\sim} \neg A, \Delta$. By the clauses of implication-space semantics this holds if and only if $\Gamma, A \overset{\mathcal{M}}{\mid\!\sim} \Delta$. Hence, by our induction hypothesis, $\Gamma, A \models_{\overline{TM}} \Delta$. But by the semantic clauses of truth-maker theory, this holds just in case $\Gamma \models_{\overline{TM}} \neg A, \Delta$. The case of $\Gamma, \neg A \overset{\mathcal{M}}{\mid\!\sim} \Delta$ is analogous.

Conjunction: Suppose that $\Gamma \overset{\mathcal{M}}{\mid\!\sim} A \wedge B, \Delta$. By the clauses of implication-space semantics this holds if and only if $\Gamma \overset{\mathcal{M}}{\mid\!\sim} A, \Delta$ and $\Gamma \overset{\mathcal{M}}{\mid\!\sim} B, \Delta$ and $\Gamma \overset{\mathcal{M}}{\mid\!\sim} A, B, \Delta$. Hence, by our induction hypothesis, $\Gamma \models_{\overline{TM}} A, \Delta$ and $\Gamma \models_{\overline{TM}} B, \Delta$ and $\Gamma \models_{\overline{TM}} A, B, \Delta$. But by the semantic clauses of truth-maker theory, this holds just in case $\Gamma \models_{\overline{TM}} A \wedge B, \Delta$.

Suppose that $\Gamma, A \wedge B \overset{\mathcal{M}}{\mid\!\sim} \Delta$. By the clauses of implication-space semantics this holds if and only if $\Gamma, A, B \overset{\mathcal{M}}{\mid\!\sim} \Delta$. Hence, by our induction hypothesis, $\Gamma, A, B \models_{\overline{TM}} \Delta$. But by the semantic clauses of truth-maker theory, this holds just in case $\Gamma, A \wedge B \models_{\overline{TM}} \Delta$. ■

Proposition 113. $\Gamma \overset{b}{\mid\!\sim} \Delta$ *just in case* $\Gamma \models_{\overline{TM}}^{\mathfrak{B}} \Delta$.

Proof. We prove a more general result, namely that if we restrict ourselves, in both theories, to models in which a given set of implications hold, then the implications that hold in all these models, in the two theories, coincide. So, let $x \subseteq \mathcal{P}(\mathfrak{L}) \times \mathcal{P}(\mathfrak{L})$, and consider only implication-space models and truth-maker models such that, for all $\langle \Gamma, \Delta \rangle \in x$, $\Gamma \overset{x}{\mid\!\sim} \Delta$ and $\Gamma \models_{\overline{TM}}^{x} \Delta$.

The consequence relations of every implication-space model and every truth-maker model are closed under the rules of NMMS and the inverted rules, which we collectively call *RUL*, because the semantic clauses of these theories are equivalent to *RUL* (as we have seen above). And if the consequence relation of every model in a set is closed under a set of rules, then the intersection of these consequence relations is closed under these rules. Hence, the consequence relations defined by any sets of models in both theories are closed under *RUL*. Let $RUL(x)$ be the closure of a set of sequents x under *RUL*. So, for all $\langle \Gamma, \Delta \rangle \in RUL(x)$, $\Gamma \overset{x}{\mid\!\sim} \Delta$ and $\Gamma \models_{\overline{TM}}^{x} \Delta$.

For the other direction, suppose that $\langle \Gamma, \Delta \rangle \notin RUL(x)$. Then we can construct counterexamples to $\Gamma \overset{x}{\mid\!\sim} \Delta$ and $\Gamma \models_{\overline{TM}}^{x} \Delta$, and we can construct these countermodels in a way that ensures that all the implications in x hold in them. In fact, the model that includes all and only the implications in $RUL(x)$ is a counterexample. Therefore, $\Gamma \overset{x}{\mid\!\sim} \Delta$ if and only if $\Gamma \models_{\overline{TM}}^{x} \Delta$.

Since restricting our attention to models in which a base consequence relation holds is a special case of restricting our attention to models in

which a particular set of implications hold, it follows that $\Gamma \overset{b}{\mid\!\sim} \Delta$ just in case $\Gamma \models^{\mathfrak{B}}_{TM} \Delta$. ■

5.7.4 Implication-Space Semantics for the Noncontractive Variant of NMMS

We now turn to the noncontractive variant of our sequent calculus, namely $\mathsf{NMMS}^{/ctr}_{\mathfrak{B}}$.

Theorem 114. *Let* $\mathfrak{B} = \left\langle \mathfrak{L}_{\mathfrak{B}}, \mathrel{\mid\!\sim}_{\mathfrak{B}} \right\rangle$ *be a base vocabulary in which contraction may fail, let* b *be the set of models that are fit for* $\mathfrak{B}$*, and let* $\Gamma, \Delta \subseteq \mathfrak{L}$*. Then* $\Gamma \overset{b}{\mid\!\sim}_{/ctr} \Delta$ *if and only if* $\Gamma \succ \Delta$ *is derivable in* $\mathsf{NMMS}^{/ctr}_{\mathfrak{B}}$.

Proof. Soundness: If $\Gamma \succ \Delta$ is derivable in $\mathsf{NMMS}^{/ctr}_{\mathfrak{B}}$, then $\Gamma \overset{b}{\mid\!\sim}_{/ctr} \Delta$. We argue by induction on the proof height of the derivation of $\Gamma \succ \Delta$. Since all models in b are fit for $\mathfrak{B}$, the base case is immediate. For the induction step, $\Gamma \succ \Delta$ comes by one of the operational rules of $\mathsf{NMMS}^{/ctr}_{\mathfrak{B}}$. Since the other connectives can be defined in terms of negation and the conditional in $\mathsf{NMMS}^{/ctr}_{\mathfrak{B}}$, it suffices to consider the rules for negation and the conditional.

Suppose $\Gamma \succ \Delta$ is $\Gamma \succ \neg A, \Delta'$ and comes by $\Gamma, A \succ \Delta'$. By our induction hypothesis, $\Gamma, A \overset{b}{\mid\!\sim}_{/ctr} \Delta'$. So for all $\mathcal{M} \in b$, labeling $\mathtt{G} = [\![\Gamma]\!]^{\mathcal{M}} = \{\mathtt{g}_0, ..., \mathtt{g}_n\}$ and $\mathtt{D}' = [\![\Delta']\!]^{\mathcal{M}} = \{\mathtt{d}_0, ..., \mathtt{d}_m\}$ and $\mathtt{a} = [\![A]\!]^{\mathcal{M}}$, we have $\mathtt{G}, \mathtt{a} \Vdash\!\sim \mathtt{D}$, which means that $\bigcup(\mathtt{a}^+ \sqcup \bigsqcup_{i=0}^{n} \mathtt{g}_i^+ \sqcup \bigsqcup_{j=0}^{m} \mathtt{d}_j^-) \subseteq \mathbb{I}$. Now, by our semantic clauses, $\mathtt{a}^+ = (\neg\mathtt{a})^-$. Hence, $\bigcup((\neg\mathtt{a})^- \sqcup \bigsqcup_{i=0}^{n} \mathtt{g}_i^+ \sqcup \bigsqcup_{j=0}^{m} \mathtt{d}_j^-) \subseteq \mathbb{I}$ and, so, $\mathtt{G} \Vdash\!\sim \neg\mathtt{a}, \mathtt{D}$. Therefore, $\Gamma \overset{b}{\mid\!\sim}_{/ctr} \neg A, \Delta'$, as desired. The case for the negation left-rule is analogous.

Suppose that $\Gamma \succ \Delta$ is $\Gamma \succ A \mathrel{\dot{\to}} B, \Delta'$ and comes by $\Gamma, A \succ B, \Delta'$. By our induction hypothesis, $\Gamma, A \overset{b}{\mid\!\sim}_{/ctr} B, \Delta'$. So for all $\mathcal{M} \in b$, labeling as before and $\mathtt{b} = [\![B]\!]^{\mathcal{M}}$, we have $\mathtt{G}, \mathtt{a} \Vdash\!\sim \mathtt{b}, \mathtt{D}$, which means that $\bigcup(\mathtt{a}^+ \sqcup \mathtt{b}^- \sqcup \bigsqcup_{i=0}^{n} \mathtt{g}_i^+ \sqcup \bigsqcup_{j=0}^{m} \mathtt{d}_j^-) \subseteq \mathbb{I}$. Now, by our semantic clauses, $\mathtt{a}^+ \sqcup \mathtt{b}^- = (\mathtt{a} \mathrel{\dot{\to}} \mathtt{b})^-$. Hence, $\bigcup((\mathtt{a} \mathrel{\dot{\to}} \mathtt{b})^- \sqcup \bigsqcup_{i=0}^{n} \mathtt{g}_i^+ \sqcup \bigsqcup_{j=0}^{m} \mathtt{d}_j^-) \subseteq \mathbb{I}$ and, so, $\mathtt{G} \Vdash\!\sim \mathtt{a} \mathrel{\dot{\to}} \mathtt{b}, \mathtt{D}$. Therefore, $\Gamma \overset{b}{\mid\!\sim}_{/ctr} A \mathrel{\dot{\to}} B, \Delta'$, as desired.

Suppose that $\Gamma \succ \Delta$ is $\Gamma', A \mathrel{\dot{\to}} B \succ \Delta$ and comes by $\Gamma' \succ A, \Delta$ and $\Gamma', B \succ \Delta$. By our induction hypothesis, $\Gamma' \overset{b}{\mid\!\sim}_{/ctr} A, \Delta$ and $\Gamma', B \overset{b}{\mid\!\sim}_{/ctr} \Delta$. By reasoning that is analogous to the reasoning in the previous two cases, this implies

that $\bigcup(\mathtt{a}^- \sqcup \bigsqcup_{i=0}^{n} \mathtt{g}_i^+ \sqcup \bigsqcup_{j=0}^{m} \mathtt{d}_j^-) \subseteq \mathbb{I}$ and $\bigcup(\mathtt{b}^+ \sqcup \bigsqcup_{i=0}^{n} \mathtt{g}_i^+ \sqcup \bigsqcup_{j=0}^{m} \mathtt{d}_j^-) \subseteq \mathbb{I}$. By our semantic clauses, $\mathtt{a}^- \sqcap \mathtt{b}^+ = (\mathtt{a} \rightarrowtriangle \mathtt{b})^+$. But, for any $\mathtt{x}$, we have $\bigcup(\mathtt{a}^- \sqcap \mathtt{b}^+ \sqcup \mathtt{x}) \subseteq \mathbb{I}$ just in case $\bigcup(\mathtt{a}^- \sqcup \mathtt{x}) \subseteq \mathbb{I}$ and $\bigcup(\mathtt{b}^+ \sqcup \mathtt{x}) \subseteq \mathbb{I}$ (by Lemma 107). So, $\mathtt{G}, \mathtt{a} \rightarrowtriangle \mathtt{b} \Vdash\!\sim \mathtt{D}$. Therefore, $\Gamma, A \rightarrowtriangle B \overset{b}{\Vdash\!\sim}_{/ctr} \Delta'$, as desired.

Completeness: Suppose that $\Gamma \succ \Delta$ is not derivable in $\mathsf{NMMS}_{\mathfrak{B}}^{/ctr}$. Then a proof-search on $\Gamma \succ \Delta$ yields an atomic sequent that is not in $\mathfrak{B}$. We can choose a model, $\mathcal{M}' \in b$, in which these atomic sequents fail, so that if $\Theta_0 \succ \Pi_0$ is an atomic sequent that results from the proof-search but is not in $\mathfrak{B}$, then $\Theta_0 \not\overset{\mathcal{M}'}{\Vdash\!\sim}_{/ctr} \Pi_0$. The latter implies that $\Gamma \not\overset{\mathcal{M}'}{\Vdash\!\sim}_{/ctr} \Delta$. For, we can show that if there is a rule-application of $\mathsf{NMMS}_{\mathfrak{B}}^{/ctr}$ that concludes some sequent $\Gamma' \succ \Delta'$ and one of the top sequents of this rule-application does not hold in a model, then $\Gamma' \succ \Delta'$ does not hold in that model. For example, if $\Gamma, A \not\overset{\mathcal{M}}{\Vdash\!\sim}_{/ctr} \Delta'$, then $\Gamma \not\overset{\mathcal{M}}{\Vdash\!\sim}_{/ctr} \neg A, \Delta'$; for if $\bigcup(\mathtt{a}^+ \sqcup \bigsqcup_{i=0}^{n} \mathtt{g}_i^+ \sqcup \bigsqcup_{j=0}^{m} \mathtt{d}_j^-) \not\subseteq \mathbb{I}$, then $\bigcup((\neg\mathtt{a})^- \sqcup \bigsqcup_{i=0}^{n} \mathtt{g}_i^+ \sqcup \bigsqcup_{j=0}^{m} \mathtt{d}_j^-) \not\subseteq \mathbb{I}$. And analogous reasoning applies to applications of the other rules. Therefore, $\Gamma \overset{b}{\not\Vdash\!\sim}_{/ctr} \Delta$, as desired. ■

5.7.5 Implication-Space Semantics for Multiplicative Additive Linear Logic

We now turn to multiplicative additive linear logic (MALL), and we start with a formulation of MALL.

Definition 115 (Phase space)**.** A phase space $\boldsymbol{P}$ consists of (i) a commutative monoid, which is a set $\boldsymbol{P}$ and an operation $\bullet$ defined on its elements, such that $\forall p, q, r \in \boldsymbol{P}$ $(p \bullet 1 = 1 \bullet p = p$ and $p \bullet q = q \bullet p$ and $(p \bullet q) \bullet r = p \bullet (q \bullet r))$; and (ii) a set of anti-phases $\perp_{\boldsymbol{P}}$.

Definition 116 (Dual)**.** If $G \subseteq \boldsymbol{P}$, then its dual $G^\perp$ is $\{p \in \boldsymbol{P} \mid \forall q \in G\,(p \bullet q \in \perp_{\boldsymbol{P}})\}$.

Definition 117 (Fact)**.** A *fact* is a subset G of $\boldsymbol{P}$ such that $G^{\perp\perp} = G$; the elements of G are called the phases of G; G is valid when $1 \in G$.

Definition 118 (Connectives)**.** $\neg G =_{df} G^\perp$; $G \otimes H =_{df} (G \bullet H)^{\perp\perp}$; $G \parr H =_{df} (G^\perp \bullet H^\perp)^\perp$; $G \& H =_{df} G \cap H$; $G \oplus H =_{df} (G \cup H)^{\perp\perp}$.

Definition 119. (i) A *phase* structure, S, for the propositional language with the connectives above consists in a phase space $\langle \boldsymbol{P}, \perp_{\boldsymbol{P}} \rangle$ and, for each propositional letter a, a fact a_S of $\boldsymbol{P}$.
(ii) With each proposition we associate its interpretation, i.e., a fact in $\boldsymbol{P}$; the interpretation of A in the structure S is denoted by A_S or $S(A)$.

(iii) A is valid in S if and only if $1 \in S(A)$,
(iv) A is a linear tautology if and only if A is valid in any phase structure S.

Definition 120 (MALL Consequence ($\models_{LL}$)). $A_1, ..., A_n \models_{LL} B_1, ..., B_m$ if and only if $A_1^{\perp} \parr ... \parr A_n^{\perp} \parr B_1 \parr ... \parr B_m$ is a linear tautology.

A couple of lemmas turn out to be useful.

Lemma 121. *There is a mapping, f, between implication frames and phase spaces such that* $\mathtt{e} \in \mathtt{RSR}(G)$ *if and only if* $1 \in (f(G))^{\perp}$.

Proof. The following is such a mapping, letting $G \circ H =_{df} \{\langle X \cup Y, Z \cup W\rangle \mid \langle X, Z\rangle \in G \ \langle Y, W\rangle \in H\}$:

1. $f(\mathbb{S}) = \boldsymbol{P}$
2. $f(G \circ H) = f(G) \bullet f(H)$
3. $f(\mathtt{e}) = 1$
4. $f(\perp_{\boldsymbol{P}}) = \mathbb{I}_{\mathcal{M}}$

The image of an implication frame under f satisfies the conditions on phase structures, as it is a commutative monoid with a set of anti-phases. Under this mapping: $f(\mathtt{RSR}(\cdot)) = (f(\cdot))^{\perp}$. For $\mathtt{RSR}(G)$ is defined as $\{x \in \mathbb{S} \mid \forall y \in G\,(x \circ y \in \mathbb{I})\}$, and the image of this set under f is $\{x \in P \mid \forall y \in f(G)\,(x \bullet y \in \perp)\}$, which is the definition of $(f(G))^{\perp}$. Hence, $\mathtt{e} \in \mathtt{RSR}(G)$ iff $f(\mathtt{e}) \in f(\mathtt{RSR}(G))$ iff $1 \in (f(G))^{\perp}$. ■

Corollary 122. $\mathtt{e} \in \mathtt{RSR}(G \circ H)$ *iff* $1 \in (f(G) \bullet f(H))^{\perp}$.

Lemma 123. $\mathtt{RSR}(G \circ H) = \mathtt{RSR}(\mathtt{g} \sqcup \mathtt{h})$.

Proof. By the definition of adjunction, $\mathtt{g} \sqcup \mathtt{h}$, is $\mathcal{R}(G \circ H)$. The role $\mathcal{R}(G \circ H)$ is the class of things whose range of subjunctive robustness is $\mathtt{RSR}(G \circ H)$. So $x \in \mathtt{g} \sqcup \mathtt{h}$ iff $\mathtt{RSR}(G \circ H) = \mathtt{RSR}(x)$. Since the range of subjunctive robustness of a set is the intersection of the ranges of subjunctive robustness of its members, $\mathtt{RSR}(\mathtt{g} \sqcup \mathtt{h})$ is $\bigcap\{\mathtt{RSR}(x) \mid x \in \mathtt{g} \sqcup \mathtt{h}\}$, which is just $\mathtt{RSR}(G \circ H)$. Therefore, $\mathtt{RSR}(G \circ H) = \mathtt{RSR}(\mathtt{g} \sqcup \mathtt{h})$. ■

Corollary 124. $\mathtt{e} \in \mathtt{RSR}(\mathtt{g} \sqcup \mathtt{h})$ *if and only if* $1 \in (f(G) \bullet f(H))^{\perp}$.

Lemma 125. *In the canonical LL models, for all sentences A such that* $[\![A]\!] = \mathtt{a}$, $\mathtt{RSR}(\mathtt{a}^{+}) = \mathtt{a}^{-}$ *and* $\mathtt{RSR}(\mathtt{a}^{-}) = \mathtt{a}^{+}$.

Proof. In the canonical LL models, all good implications can be represented in the form $\bigcup(\mathtt{a}^+ \sqcup \mathtt{a}^-) \subseteq \mathbb{I}$. So if $\bigcup(\mathtt{a}^+ \sqcup \mathtt{x}) \subseteq \mathbb{I}$, then $\mathtt{x} = \mathtt{a}^-$. Hence, $\mathtt{RSR}(\mathtt{a}^+) = \mathtt{a}^-$. And the analogous reasoning works for $\mathtt{RSR}(\mathtt{a}^-) = \mathtt{a}^+$. ■

Lemma 126. *For all phase structures S and the canonical LL models that are connected by a mapping like f above: If* $1 \in A_S$ *iff* $\mathtt{e} \in \mathtt{RSR}(\mathtt{a}^-)$ *and* $\mathtt{e} \in \mathtt{RSR}(\mathtt{a}^+)$ *iff* $1 \in A_S^\perp$ *both hold for the sentences and roles mentioned in the following statements, then these statements hold (with square brackets indicating alternatives that also hold, when uniformly applied in a given statement).*

(a) $\mathtt{e} \in \mathtt{RSR}(\mathtt{a}^{+[/-]} \sqcup \mathtt{b}^{+[/-]})$ *iff* $1 \in (A_S^{[\perp]} \bullet B_S^{[\perp]})^\perp$

(b) $\mathtt{e} \in \mathtt{RSR}(\mathtt{a}^{-[/+]} \sqcap \mathtt{b}^{-[/+]})$ *iff* $1 \in (A_S^{[\perp]} \cap B_S^{[\perp]})$

(c) $\mathtt{e} \in \mathtt{RSR}(\mathtt{RSR}(\mathtt{a}^{+[/-]} \cup \mathtt{b}^{+[/-]}))$ *iff* $1 \in (A_S^{[\perp]} \cup B_S^{[\perp]})^{\perp\perp}$

Proof. (ad a) The first fact is immediate from Corollary 124 and Lemma 125.

(ad b) $\mathtt{e} \in \mathtt{RSR}(\mathtt{a}^- \sqcap \mathtt{b}^-)$ iff $\mathtt{e} \in \mathtt{RSR}(\mathtt{a}^-) \cap \mathtt{RSR}(\mathtt{b}^-)$ iff $1 \in (A_S \cap B_S)$. And Lemma 125, ensures that $\mathtt{e} \in \mathtt{RSR}(\mathtt{a}^+ \sqcap \mathtt{b}^+)$ iff $1 \in (A_S^\perp \cap B_S^\perp)$.

(ad c) $\mathtt{e} \in \mathtt{RSR}(\mathtt{RSR}(\mathtt{a}^+ \cup \mathtt{b}^+))$ iff $\mathtt{e} \in \mathtt{RSR}(\mathtt{RSR}(\mathtt{a}^+) \cap \mathtt{RSR}(\mathtt{b}^+))$ iff $\mathtt{e} \in \mathtt{RSR}(\mathtt{a}^-) \cup \mathtt{RSR}(\mathtt{b}^-)$ iff $1 \in (A_S \cup B_S)$ iff $1 \in (A_S \cup B_S)^{\perp\perp}$. ■

Proposition 127. *In the canonical LL models and for all phase structures, letting* $f(\mathtt{a}) = A_S$ *for all roles and facts,* $1 \in A_S$ *iff* $\mathtt{e} \in \mathtt{RSR}(\mathtt{a}^-)$ *and* $\mathtt{e} \in \mathtt{RSR}(\mathtt{a}^+)$ *iff* $1 \in A_S^\perp$.

Proof. We argue by induction on the complexity of A. The base case is trivial because no atomic sentences are linear tautologies, and if A is an atomic sentence then $\mathtt{e} \notin \mathtt{RSR}(\langle \emptyset, p \rangle)$ in the canonical LL models. And similarly for $A_S^\perp$ and $\mathtt{RSR}(\langle p, \emptyset \rangle)$.

For the induction step, let's start with negation: $\mathtt{e} \in \mathtt{RSR}((\neg\mathtt{a})^-)$ iff $\mathtt{e} \in \mathtt{RSR}(\mathtt{a}^+)$ iff $1 \in A_S^\perp$ iff $1 \in (\neg A)_S$. Similarly, $\mathtt{e} \in \mathtt{RSR}((\neg\mathtt{a})^+)$ iff $\mathtt{e} \in \mathtt{RSR}(\mathtt{a}^-)$ iff $1 \in A_S$ iff $1 \in (\neg A)_S^\perp$.

Multiplicative conjunction: $\mathtt{e} \in \mathtt{RSR}((\mathtt{a} \otimes \mathtt{b})^-)$ iff $\mathtt{e} \in \mathtt{RSR}(\mathtt{RSR}(\mathtt{RSR}(\mathtt{a}^-) \sqcup \mathtt{RSR}(\mathtt{b}^-)))$ iff $\mathtt{e} \in \mathtt{RSR}(\mathtt{RSR}(\mathtt{a}^+ \sqcup \mathtt{b}^+))$ iff $1 \in (A_S \bullet B_S)^{\perp\perp}$ (by Lemma 126(a)) iff $1 \in (A \otimes B)_S$. Similarly, $\mathtt{e} \in \mathtt{RSR}((\mathtt{a} \otimes \mathtt{b})^+)$ iff $\mathtt{e} \in \mathtt{RSR}(\mathtt{a}^+ \sqcup \mathtt{b}^+)$ iff $1 \in (A_S \bullet B_S)^\perp$ (by Lemma 126(a)) iff $1 \in (A_S \bullet B_S)^{\perp\perp\perp}$ iff $1 \in (A \otimes B)_S^\perp$.

Multiplicative disjunction: $\mathtt{e} \in \mathtt{RSR}((\mathtt{a} \parr \mathtt{b})^-)$ iff $\mathtt{e} \in \mathtt{RSR}(\mathtt{a}^- \sqcup \mathtt{b}^-)$ iff $1 \in (A_S^\perp \bullet B_S^\perp)^\perp$ (by Lemma 126(a)) iff $1 \in (A \parr B)_S$. Similarly $\mathtt{e} \in \mathtt{RSR}((\mathtt{a} \parr \mathtt{b})^+)$ iff $\mathtt{e} \in \mathtt{RSR}(\mathtt{RSR}(\mathtt{RSR}(\mathtt{a}^+) \sqcup \mathtt{RSR}(\mathtt{b}^+)))$ iff $\mathtt{e} \in \mathtt{RSR}(\mathtt{RSR}(\mathtt{a}^- \sqcup \mathtt{b}^-))$ iff $1 \in (A_S^\perp \bullet B_S^\perp)^{\perp\perp}$ (by Lemma 126(a)) iff $1 \in (A \parr B)_S^\perp$.

Additive conjunction: $\mathtt{e} \in \mathtt{RSR}((\mathtt{a\&b})^-)$ iff $\mathtt{e} \in \mathtt{RSR}(\mathtt{a}^- \sqcap \mathtt{b}^-)$ iff $1 \in (A_S \cap B_S)$ (by Lemma 126(b)) iff $1 \in (A\&B)_S$. Similarly, $\mathtt{e} \in \mathtt{RSR}((\mathtt{a\&b})^+)$ iff $\mathtt{e} \in \mathtt{RSR}(\mathtt{RSR}(\mathtt{RSR}(\mathtt{a}^+) \cup \mathtt{RSR}(\mathtt{b}^+)))$ iff $\mathtt{e} \in \mathtt{RSR}(\mathtt{RSR}(\mathtt{a}^- \cup \mathtt{b}^-))$ iff $1 \in (A_S^\perp \cup B_S^\perp)^{\perp\perp}$ (by Lemma 126(c)) iff $1 \in (\neg A \oplus \neg B)_S$ iff $1 \in (A\&B)_S^\perp$.

Additive disjunction: $\mathtt{e} \in \mathtt{RSR}((\mathtt{a} \oplus \mathtt{b})^-)$ iff $\mathtt{e} \in \mathtt{RSR}(\mathtt{RSR}(\mathtt{RSR}(\mathtt{a}^-) \cup \mathtt{RSR}(\mathtt{b}^-)))$ iff $\mathtt{e} \in \mathtt{RSR}(\mathtt{RSR}(\mathtt{a}^+ \cup \mathtt{b}^+))$ iff $1 \in (A_S \cup B_S)^{\perp\perp}$ (by Lemma 126(c)) iff $1 \in (A \oplus B)_S$. Similarly, $\mathtt{e} \in \mathtt{RSR}((\mathtt{a} \oplus \mathtt{b})^+)$ iff $\mathtt{e} \in \mathtt{RSR}(\mathtt{a}^+ \sqcap \mathtt{b}^+)$ iff $1 \in (A_S^\perp \cap B_S^\perp)$ (by Lemma 126(b)) iff $1 \in (\neg A \& \neg B)_S$ iff $1 \in (A \oplus B)_S^\perp$. ■

Theorem 128. *A is a linear tautology if and only if $\overset{\mathcal{M}}{\vdash\!\!\sim} A$ in all LL implication-space models. And, hence, $\Gamma \models_{LL} \Delta$ iff $\Gamma \overset{LL}{\vdash\!\!\sim} \Delta$.*

Proof. By Proposition 127, $1 \in A_S$ in all phase structures, which means that A is a linear tautology, if and only if $\mathtt{e} \in \mathtt{RSR}(\mathtt{a}^-)$ and, hence, $\overset{\mathcal{M}}{\vdash\!\!\sim} A$ in the canonical LL implication-space models. So, A is a linear tautology if and only if $\overset{\mathcal{M}}{\vdash\!\!\sim} A$ in all LL implication-space models. Let $\Gamma = \{\gamma_1, ..., \gamma_n\}$ and $\Delta = \{\delta_1, ..., \delta_n\}$. Then $\Gamma \models_{LL} \Delta$ iff $\gamma_1^\perp \parr ... \parr \gamma_n^\perp \parr \delta_1 \parr ... \parr \delta_m$ is a linear tautology. And the latter holds iff $\overset{\mathcal{M}}{\vdash\!\!\sim} \gamma_1^\perp \parr ... \parr \gamma_n^\perp \parr \delta_1 \parr ... \parr \delta_m$ in all LL implication-space models, which holds just in case $\Gamma \overset{LL}{\vdash\!\!\sim} \Delta$. ■

5.7.6 Implicational Role Inclusions

We now turn to implicational role inclusions. However, in order to give the proofs of our results regarding role inclusions, we must first establish some connections between implication-space models and strong Kleene valuations.

Definition 129 (Base Models of a Language $\mathfrak{L}$)**.** Let $\mathfrak{L}_0$ be the set of atomic sentences in $\mathfrak{L}$. A base model of the language $\mathfrak{L}$ is a pair of contents defined by the implication frame, $\langle \mathfrak{L}_0, \mathbb{I} \rangle$, where the interpretation is such that for all $\phi \in \mathfrak{L}$, $[\![\phi]\!] = \langle \mathcal{R}\langle\{\phi\}, \emptyset\rangle, \mathcal{R}\langle\emptyset, \{\phi\}\rangle\rangle$. Base models are monotonic just in case $\forall \mathcal{R}_1, \mathcal{R}_2 \in \mathbb{R}$ (If $\bigcup(\mathcal{R}_1) \subseteq \mathbb{I}$, then $\bigcup((\mathcal{R}_1 \sqcup \mathcal{R}_2)) \subseteq \mathbb{I}$). And they are reflexive just in case $\forall [\![\phi]\!] \in \mathbb{C}\ \forall \mathcal{R} \in \mathbb{R}\ \bigcup(([\![\phi]\!]^+ \sqcup [\![\phi]\!]^- \sqcup \mathcal{R})) \subseteq \mathbb{I}$.

Definition 130 (Strong Kleene valuation (model, interpretation))**.** A strong Kleene valuation is a function from sentences of a language $\mathfrak{L}$ to truth-values in $\{0, \frac{1}{2}, 1\}$ that obeys the strong Kleene truth-tables which are, for negation and conjunction (with disjunction and conditionals being defined in the familiar way):

- $v(\neg\phi) = 1 - v(\phi)$
- $v(\phi \wedge \psi) = min(v(\phi), v(\psi))$

Definition 131 (Maximization of an atomic implication)**.** Let $\langle \Gamma_0, \Delta_0 \rangle$ be a pair of sets of atomic bearers. Then the maximization of it, written $Mx\,\langle \Gamma_0, \Delta_0 \rangle$, is the pair of bearers $\langle \Gamma, \Delta \rangle = \langle \bigcup \Gamma_i, \bigcup \Delta_i \rangle$ for all $\langle \Gamma_i, \Delta_i \rangle$ in the series from $\langle \Gamma_0, \Delta_0 \rangle$ to $\langle \Gamma_\omega, \Delta_\omega \rangle$ inductively defined as follows:

(0) $\Gamma_n \subseteq \Gamma_{n+1}$ and $\Delta_n \subseteq \Delta_{n+1}$.
(i) If $\phi \in \Gamma_n$, then $\neg\phi \in \Delta_{n+1}$. And if $\phi \in \Delta_n$, then $\neg\phi \in \Gamma_{n+1}$.
(ii) If $\phi \in \Gamma_n$ and $\psi \in \Gamma_n$, then $\phi \wedge \psi \in \Gamma_{n+1}$.
(iii) If $\phi \in \Delta_n$ or $\psi \in \Delta_n$, then $\phi \wedge \psi \in \Delta_{n+1}$.

Lemma 132. *There is a strong Kleene valuation, v, that partitions all sentences of language $\mathfrak{L}$ such that $v(\gamma) = 1$ iff $\gamma \in \Gamma$, and $v(\delta) = 0$ iff $\delta \in \Delta$, and $v(\theta) = \frac{1}{2}$ iff $\theta \in \Theta$, if and only if there is a monotonic and reflexive base model of $\mathfrak{L}$, such that $\mathbb{I} = \mathcal{P}(\mathfrak{L}_0) \times \mathcal{P}(\mathfrak{L}_0) \setminus \{\langle X, Y \rangle \mid X \subseteq \Gamma_0, Y \subseteq \Delta_0\}$ and $\langle \Gamma, \Delta \rangle = Mx\,\langle \Gamma_0, \Delta_0 \rangle$.*

Proof. Left-to-right: Suppose that there is a strong Kleene valuation that partitions all sentences such that $v(\gamma) = 1$ iff $\gamma \in \Gamma$, and $v(\delta) = 0$ iff $\delta \in \Delta$, and $v(\theta) = \frac{1}{2}$ iff $\theta \in \Theta$. Let $\Gamma_0 = \Gamma \cap \mathfrak{L}_0$ be the atomic sentences in Γ, and let $\Delta_0 = \Delta \cap \mathfrak{L}_0$ be the atomic sentences in Δ. The base model in which $\mathbb{I} = \mathcal{P}(\mathfrak{L}_0) \times \mathcal{P}(\mathfrak{L}_0) \setminus \{\langle X, Y \rangle \mid X \subseteq \Gamma_0, Y \subseteq \Delta_0\}$ is the desired model. To see this, note that since $\Gamma_0 \cap \Delta_0 = \emptyset$, the stipulation that $\{\langle X, Y \rangle \mid X \subseteq \Gamma_0, Y \subseteq \Delta_0\} \not\subseteq \mathbb{I}$ does not conflict with the reflexivity of the base model. Moreover, suppose that $\bigcup(\mathcal{R}_1) \subseteq \mathbb{I}$. So, $\bigcup(\mathcal{R}_1) \cap \{\langle X, Y \rangle \mid X \subseteq \Gamma_0, Y \subseteq \Delta_0\} = \emptyset$. But then for any $\mathcal{R}_2$ the adjunction $\bigcup(\mathcal{R}_1 \sqcup \mathcal{R}_2) \cap \{\langle X, Y \rangle \mid X \subseteq \Gamma_0, Y \subseteq \Delta_0\} = \emptyset$. So, $\bigcup(\mathcal{R}_1 \sqcup \mathcal{R}_2) \subseteq \mathbb{I}$. Hence, $\forall \mathcal{R}_1, \mathcal{R}_2 \in \mathbb{R}$ (if $\bigcup(\mathcal{R}_1) \subseteq \mathbb{I}$, then $\bigcup(\mathcal{R}_1 \sqcup \mathcal{R}_2) \subseteq \mathbb{I}$), and so our stipulation implies that our base model is monotonic.

To show that $\langle \Gamma, \Delta \rangle = Mx\,\langle \Gamma_0, \Delta_0 \rangle$, we argue by induction on steps in our construction of $Mx\,\langle \Gamma_0, \Delta_0 \rangle$, namely that for all sentences, ϕ_n, of complexity n, we have $v(\phi_n) = 1$ iff $\phi_n \in \Gamma_n$, and $v(\phi_n) = 0$ iff $\phi_n \in \Delta_n$. All the other sentences of any complexity must then be in Θ. Base case: For $\langle \Gamma_0, \Delta_0 \rangle$ it is immediate that it is the pair of atomic sentences, i.e., sentences of complexity 0, from Γ and Δ. So that, $v(\phi_0) = 1$ iff $\phi_0 \in \Gamma_0$, and $v(\phi_0) = 0$ iff $\phi_0 \in \Delta_0$.

For the induction step, take a sentence, ϕ, of complexity, $n + 1$. If ϕ_{n+1} is either in Γ or in Δ, then it is added at step $n + 1$ of the construction of $Mx\,\langle \Gamma_0, \Delta_0 \rangle$. And ϕ is either a negation or a conjunction. Suppose $\phi = \neg\psi$. Now, $\neg\psi \in \Gamma_{n+1}$ iff $\psi \in \Delta_n$, which happens, by our induction hypothesis iff $v(\psi) = 0$, which by the truth-tables happens iff $v(\neg\psi) = 1$. Similarly, $\neg\psi \in \Delta_{n+1}$ iff $\psi \in \Gamma_n$, which happens, by our induction hypothesis iff

$v(\psi) = 1$, which by the truth-tables happens iff $v(\neg\psi) = 0$. Next, suppose that $\phi = \psi \wedge \chi$. Now, $\psi \wedge \chi \in \Gamma_{n+1}$ iff $\psi \in \Gamma_n$ and $\chi \in \Gamma_n$, which happens, by our induction hypothesis iff $v(\psi) = 1$ and $v(\chi) = 1$, which by the truth-tables happens iff $v(\psi \wedge \chi) = 1$. Similarly, $\psi \wedge \chi \in \Delta_{n+1}$ iff $\psi \in \Delta_n$ or $\chi \in \Delta_n$, which happens, by our induction hypothesis iff $v(\psi) = 0$ or $v(\chi) = 0$, which by the truth-tables happens iff $v(\psi \wedge \chi) = 0$. So, by induction, $\langle \Gamma, \Delta \rangle = Mx \langle \Gamma_0, \Delta_0 \rangle$.

Right-to-left: Suppose that there is a monotonic and reflexive base model of $\mathfrak{L}$, such that $\mathbb{I} = \mathcal{P}(\mathfrak{L}_0) \times \mathcal{P}(\mathfrak{L}_0) \setminus \{\langle X, Y \rangle \mid X \subseteq \Gamma_0, Y \subseteq \Delta_0\}$ and $\langle \Gamma, \Delta \rangle = Mx \langle \Gamma_0, \Delta_0 \rangle$. Take the valuation such that $v(\gamma_0) = 1$ iff $\gamma_0 \in \Gamma_0$, and $v(\delta_0) = 0$ iff $\delta_0 \in \Delta_0$, and $v(\theta_0) = \frac{1}{2}$ iff $\theta_0 \in \mathfrak{L}_0$ but $\theta_0 \notin \Gamma_0 \cup \Delta_0$. Then we use the strong Kleene truth-tables to determine the values of all logically complex sentences. The result is a strong Kleene valuation, v, that partitions all sentences of language $\mathfrak{L}$ such that $v(\gamma) = 1$ iff $\gamma \in \Gamma$, and $v(\delta) = 0$ iff $\delta \in \Delta$, and $v(\theta) = \frac{1}{2}$ iff $\theta \in \Theta$. This can be shown by an induction on the steps in the construction of $Mx \langle \Gamma_0, \Delta_0 \rangle$, which is analogous to the induction for the left-to-right direction. ■

Lemma 133. *For every set $\mathscr{I} \subseteq \mathcal{P}(\mathfrak{L}_0) \times \mathcal{P}(\mathfrak{L}_0)$, there is a monotonic and reflexive base model of $\mathfrak{L}$, such that $\mathbb{I} = \mathcal{P}(\mathfrak{L}_0) \times \mathcal{P}(\mathfrak{L}_0) \setminus \{\langle X, Y \rangle \mid X \subseteq \Gamma_0, Y \subseteq \Delta_0\}$ and $\langle \Gamma, \Delta \rangle = Mx \langle \Gamma_0, \Delta_0 \rangle$ and $\mathscr{I} = \{\langle X, Y \rangle \mid X \subseteq \Gamma, Y \subseteq \Delta\}$ if and only if there is a strong Kleene valuation, v, that is an* ST *counterexample to all and only the inferences in $\mathscr{I}$.*

Proof. Left-to-right: For all candidate implications $\langle X, Y \rangle \in \mathscr{I}$, we know that $X \subseteq \Gamma$ and $Y \subseteq \Delta$. By Lemma 132, there is a strong Kleene valuation such that $v(\gamma) = 1$ iff $\gamma \in \Gamma$, and $v(\delta) = 0$ iff $\delta \in \Delta$, and $v(\theta) = \frac{1}{2}$ iff $\theta \in \Theta$. Hence, for all $\langle X, Y \rangle \in \mathscr{I}, \gamma \in X(v(\gamma) = 1)$ and $\forall \delta \in Y(v(\delta) = 0)$. Hence, v is an ST counterexample to everything in $\mathscr{I}$.

Right-to-left: Suppose that there is a strong Kleene valuation, v, that is an ST counterexample to all and only the inferences in $\mathscr{I}$. By Lemma 132, there is a monotonic and reflexive base model of $\mathfrak{L}$, such that candidate implication $\mathbb{I} = \mathcal{P}(\mathfrak{L}_0) \times \mathcal{P}(\mathfrak{L}_0) \setminus \{\langle X, Y \rangle \mid X \subseteq \Gamma_0, Y \subseteq \Delta_0\}$ and $Mx(\langle \Gamma_0, \Delta_0 \rangle) = \langle \Gamma, \Delta \rangle$ such that Γ is the exact set of sentences with value 1 in v and Δ is the exact set of the sentences with value 0 in v. Hence, the inferences to which v provides ST counterexamples are all and only the inferences, $X \succ Y$, such that $X \subseteq \Gamma$ and $Y \subseteq \Delta$. That is, $\mathscr{I} = \{\langle X, Y \rangle \mid X \subseteq \Gamma, Y \subseteq \Delta\}$. ■

According to the definition in the main text, a conic implication-space model, $\mathcal{M}$, is a monotonic and reflexive implication-space model in which there is a pair $\langle \Gamma, \Delta \rangle$ and $X \overset{\mathcal{M}}{\not\succ} Y$ iff $\langle X, Y \rangle \in \mathcal{P}(\Gamma) \times \mathcal{P}(\Delta)$.

Proposition 134. *There is a monotonic and reflexive base model of* $\mathfrak{L}$*, such that* $\mathbb{I} = \mathcal{P}(\mathfrak{L}_0) \times \mathcal{P}(\mathfrak{L}_0) \setminus \{\langle X, Y\rangle \mid X \subseteq \Gamma_0, Y \subseteq \Delta_0\}$ *and* $\langle \Gamma, \Delta\rangle = Mx\,\langle \Gamma_0, \Delta_0\rangle$ *if and only if there is a conic implication-space model that is a counterexample to all and only the candidate implications in the set* $\{\langle X, Y\rangle \mid X \subseteq \Gamma, Y \subseteq \Delta\}$.

Proof. Left-to-right: Suppose that there is a monotonic and reflexive base model of $\mathfrak{L}$, such that $\mathbb{I} = \mathcal{P}(\mathfrak{L}_0) \times \mathcal{P}(\mathfrak{L}_0) \setminus \{\langle X, Y\rangle \mid X \subseteq \Gamma_0, Y \subseteq \Delta_0\}$ and $\langle \Gamma, \Delta\rangle = Mx\,\langle \Gamma_0, \Delta_0\rangle$. We know that $\langle \Gamma_0, \Delta_0\rangle \notin \mathbb{I}$. But by the construction of $Mx\,\langle \Gamma_0, \Delta_0\rangle$, if $\langle \Gamma_i, \Delta_i\rangle \notin \mathbb{I}$, then $\langle \Gamma_{i+1}, \Delta_{i+1}\rangle \notin \mathbb{I}$. Therefore, $\langle \Gamma, \Delta\rangle \notin \mathbb{I}$. And by monotonicity, this implies that none of the implications in the set $\{\langle X, Y\rangle \mid X \subseteq \Gamma, Y \subseteq \Delta\}$ are in $\mathbb{I}$. Conversely, suppose that $\langle U, W\rangle \notin \mathbb{I}$. If the sentences in U and W are atomic, then $\langle U, W\rangle \in \{\langle X, Y\rangle \mid X \subseteq \Gamma_0, Y \subseteq \Delta_0\}$. And if they are not atomic, we can reduce them to atomic sequents by applying the procedure by which we constructed $Mx\,\langle \Gamma_0, \Delta_0\rangle$ in reverse. The resulting sequents must be in $\{\langle X, Y\rangle \mid X \subseteq \Gamma_0, Y \subseteq \Delta_0\}$. And from the construction of $Mx\,\langle \Gamma_0, \Delta_0\rangle$ we can then see that $\langle U, W\rangle \in \{\langle X, Y\rangle \mid X \subseteq \Gamma, Y \subseteq \Delta\}$. So our model is a counterexample to all and only the candidate implications in the set $\{\langle X, Y\rangle \mid X \subseteq \Gamma, Y \subseteq \Delta\}$.

Right-to-left: Suppose that there is a monotonic and reflexive implication-space model that is a counterexample to all and only the candidate implications in the set $\{\langle X, Y\rangle \mid X \subseteq \Gamma, Y \subseteq \Delta\}$. We can now construct a model, $\mathcal{M}'$, by taking the implication frame defined by the set of atomic sentences as bearers and the consequence relation of our implication-space model over these atomic sentences as the set of good implications, and it is easy to see that the consequence relation of this base model is a monotonic and reflexive base model of $\mathfrak{L}$, such that $\mathbb{I}_{\mathcal{M}'} = \mathcal{P}(\mathfrak{L}_0) \times \mathcal{P}(\mathfrak{L}_0) \setminus \{\langle X, Y\rangle \mid X \subseteq \Gamma_0, Y \subseteq \Delta_0\}$ and $\langle \Gamma, \Delta\rangle = Mx\,\langle \Gamma_0, \Delta_0\rangle$. ■

Theorem 135. *There is a strong Kleene model that is an* ST *counterexample to all and only the candidate implications in a set* $\mathscr{K}$ *if and only if there is a conic implication-space model that is a counterexample to all and only the candidate implications in a set* $\mathscr{K}$.

Proof. Left-to-right: There is a strong Kleene model, v, that is an ST counterexample to all and only the candidate implications in a set $\mathscr{K}$ iff for all and only the implications $\langle P, C\rangle \in \mathscr{K}$ we have $v(p) = 1$ for all $p \in P$, and $v(c) = 0$ for all $c \in C$. Let Θ be $\bigcup P$ for all $\langle P, C\rangle \in \mathscr{K}$, and let Λ be $\bigcup C$ for all $\langle P, C\rangle \in \mathscr{K}$. The counterexample in question is one that partitions all sentences of language $\mathfrak{L}$ such that $v(\theta) = 1$ iff $\theta \in \Theta$, and $v(\lambda) = 0$ iff $\lambda \in \Lambda$. By Lemma 132, there

is a strong Kleene valuation, v, that partitions all sentences of language $\mathfrak{L}$ such that $v(\theta) = 1$ iff $\theta \in \Theta$, and $v(\lambda) = 0$ iff $\lambda \in \Lambda$, if and only if there is a monotonic and reflexive base model of $\mathfrak{L}$, such that candidate implication $\mathbb{I} = \mathcal{P}(\mathfrak{L}_0) \times \mathcal{P}(\mathfrak{L}_0) \setminus \{\langle X, Y\rangle \mid X \subseteq \Gamma_0, Y \subseteq \Delta_0\}$ and $\langle \Gamma, \Delta\rangle = Mx\,\langle \Gamma_0, \Delta_0\rangle$. And from Lemma 133 we know that the bad candidate implications in such a model are $\mathscr{K} = \{\langle X, Y\rangle \mid X \subseteq \Gamma, Y \subseteq \Delta\}$. Hence, there is a conic implication-space model that is a counterexample to all and only the candidate implications in a set $\mathscr{K}$.

Right-to-left: Suppose that there is a conic implication-space model that is a counterexample to all and only the candidate implications in a set $\mathscr{K}$. The set of all bad candidate implications is $\mathscr{K} = \{\langle X, Y\rangle \mid X \subseteq \Gamma, Y \subseteq \Delta\}$ for some $\langle \Gamma, \Delta\rangle$. By Proposition 134, there is such a conic implication-space model iff there is a monotonic and reflexive base model of $\mathfrak{L}$, such that $\mathbb{I} = \mathcal{P}(\mathfrak{L}_0) \times \mathcal{P}(\mathfrak{L}_0) \setminus \{\langle X, Y\rangle \mid X \subseteq \Gamma_0, Y \subseteq \Delta_0\}$ and $\langle \Gamma, \Delta\rangle = Mx\,\langle \Gamma_0, \Delta_0\rangle$. And by Lemma 132, this happens just in case there is a strong Kleene model that is an ST counterexample to all and only the candidate implications in a set $\mathscr{K}$. ■

Definition 136 (Local metainferential validity)**.** A metainference from the set of sequents X to the set of sequents Y is valid in the local sense if and only if, for every model (valuation) (in the relevant class of models), if the model is a counterexample to all of the sequents in Y, then it is a counterexample to at least one of the sequents in X.

Fact 137. *The following are implications of well known facts about the relations between* ST, LP, K3, *and* TS *(see Dicher and Paoli, 2019; Barrio et al., 2015; Cobreros et al., 2020a; Ré et al., 2021):*

$\Gamma \models_{LP} \Delta$ *iff the metainference from* $\succ \bigwedge \Gamma$ *to* $\succ \Delta$ *is locally valid in the strong Kleene formulation of* ST.

$\Gamma \models_{K3} \Delta$ *iff the metainference from* $\bigvee \Delta \succ$ *to* $\Gamma \succ$ *is locally valid in the strong Kleene formulation of* ST.

$\Gamma \models_{ST} \Delta$ *iff the metainference from* $A \succ A$ *to* $\Gamma \succ \Delta$ *is locally valid in the strong Kleene formulation of* ST.

$\Gamma \models_{TS} \Delta$ *iff the metainference from* $\{\succ \bigwedge \Gamma,\ \bigvee \Delta \succ\}$ *to* $\emptyset \succ \emptyset$ *is locally valid in the strong Kleene formulation of* ST.

Proof. One way to see why these facts hold is the following: A strong Kleene counterexample to $\Gamma \models_{LP} \Delta$ is such that $\forall \gamma \in \Gamma\ (v(\gamma) \neq 0)$ and $\forall \delta \in \Delta\ (v(\delta) = 0)$. But a strong Kleene model is an ST counterexample to $\succ \bigwedge \Gamma$ iff $\exists \gamma \in \Gamma\ (v(\gamma) = 0)$. So the models such that $\forall \gamma \in \Gamma\ (v(\gamma) \neq 0)$ are those that could be counterexamples to the metainference or the LP inference, if we considered just Γ. A counterexample to $\succ \Delta$ is a model

such that $\forall\delta \in \Delta\ (v(\delta) = 0)$. So a model is a counterexample to the metainference just in case it is a counterexample to $\Gamma \models_{LP} \Delta$.

Similarly, a strong Kleene counterexample to $\Gamma \models_{K3} \Delta$ is such that $\forall\gamma \in \Gamma\ (v(\gamma) = 1)$ and $\forall\delta \in \Delta\ (v(\delta) \neq 1)$. But a strong Kleene model is an ST counterexample to $\bigvee \Delta \succ$ iff $\exists\delta \in \Delta\ (v(\delta) = 1)$. So the models such that $\forall\delta \in \Delta\ (v(\delta) \neq 1)$ are those that could be counterexamples to the metainference or the K3 inference, if we considered just Δ. A counterexample to $\Gamma \succ$ is a model such that $\forall\gamma \in \Gamma\ (v(\gamma) = 1)$. So a model is a counterexample to the metainference just in case it is a counterexample to $\Gamma \models_{K3} \Delta$.

Of course, $\Gamma \models_{ST} \Delta$ holds iff there is no ST counterexample to $\Gamma \succ \Delta$. Since there isn't any countermodel to the inference $A \succ A$ in the strong Kleene formulation of ST, this holds just in case the metainference from $A \succ A$ to $\Gamma \succ \Delta$ is locally valid in the strong Kleene formulation of ST.

A strong Kleene counterexample to $\Gamma \models_{TS} \Delta$ is a model in which $\forall\gamma \in \Gamma$ $(v(\gamma) \neq 0)$ and $\forall\delta \in \Delta\ (v(\delta) \neq 1)$. So $\Gamma \models_{TS} \Delta$ iff every model is such that $\exists\gamma \in \Gamma\ (v(\gamma) = 0)$ or $\exists\delta \in \Delta\ (v(\delta) = 1)$. That is, $\Gamma \models_{TS} \Delta$ iff every model is an ST counterexample to $\succ \bigwedge \Gamma$ or to $\bigvee \Delta \succ$. Since all strong Kleene models are counterexamples to $\emptyset \succ \emptyset$, this holds just in case the metainference from $\{\succ \bigwedge \Gamma,\ \bigvee \Delta \succ\}$ to $\emptyset \succ \emptyset$ is locally valid in the strong Kleene formulation of ST. ■

Lemma 138. *A metainference from* $\{\Theta_1 \succ \Lambda_1, ..., \Theta_n \succ \Lambda_n\}$ *to* $\{\Gamma \succ \Delta\}$ *is locally valid in the strong Kleene formulation of* ST *just in case* $\mathcal{R}\langle\Theta_1, \Lambda_1\rangle, ..., \mathcal{R}\langle\Theta_n, \Lambda_n\rangle \preceq \mathcal{R}\langle\Gamma, \Delta\rangle$ *in the conic implication-space models.*

Proof. Left-to-right: Suppose that the metainference from $\{\Theta_1 \succ \Lambda_1, ..., \Theta_n \succ \Lambda_n\}$ to $\{\Gamma \succ \Delta\}$ is locally valid in the strong Kleene formulation of ST. By Theorem 135 above, there is a one-to-one mapping between conic implication-space models and ST models that preserves counterexamples in both directions. Hence, every conic implication-space model that is a counterexample to $\Gamma \succ \Delta$ is a counterexample to at least one of the implications in $\{\Theta_1 \succ \Lambda_1, ..., \Theta_n \succ \Lambda_n\}$. So, in all conic implication-space models, if $\mathtt{e} \notin \mathtt{RSR}(\mathcal{R}\langle\Gamma, \Delta\rangle)$, then $\mathtt{e} \notin \mathtt{RSR}(\mathcal{R}\langle\Theta_1, \Lambda_1\rangle)$ or ... or $\mathtt{e} \notin \mathtt{RSR}(\mathcal{R}\langle\Theta_n, \Lambda_n\rangle)$, i.e., $\mathtt{e} \notin \mathtt{RSR}(\mathcal{R}\langle\Theta_1, \Lambda_1\rangle \sqcap ... \sqcap \mathcal{R}\langle\Theta_n, \Lambda_n\rangle)$. In other words, it holds just in case: in all models, if $\mathtt{e} \in \mathtt{RSR}(\mathcal{R}\langle\Theta_1, \Lambda_1\rangle \sqcap ... \sqcap \mathcal{R}\langle\Theta_n, \Lambda_n\rangle)$, then $\mathtt{e} \in \mathtt{RSR}(\mathcal{R}\langle\Gamma, \Delta\rangle)$. Since we can vary $\mathbb{I}$ in any way that is consistent with the model being conic, this holds just in case, in all conic models, if $x \in \mathtt{RSR}(\mathcal{R}\langle\Theta_1, \Lambda_1\rangle \sqcap ... \sqcap \mathcal{R}\langle\Theta_n, \Lambda_n\rangle)$, then $x \in \mathtt{RSR}(\mathcal{R}\langle\Gamma, \Delta\rangle)$, that is, $\mathtt{RSR}(\bigsqcap_{i=1}^{n} \mathcal{R}\langle\Theta_i, \Lambda_i\rangle) \subseteq \mathtt{RSR}(\bigsqcup_{j=k}^{m} \mathcal{R}\langle\Gamma, \Delta\rangle)$. By definition, this is $\mathcal{R}\langle\Theta_1, \Lambda_1\rangle, ..., \mathcal{R}\langle\Theta_n, \Lambda_n\rangle \preceq \mathcal{R}\langle\Gamma, \Delta\rangle$ in the conic models.

Right-to-left: Suppose that $\mathcal{R}\langle\Theta_1,\Lambda_1\rangle,\ldots,\mathcal{R}\langle\Theta_n,\Lambda_n\rangle \preceq \mathcal{R}\langle\Gamma,\Delta\rangle$ in the conic implication-space models. So, in all conic implication-space models, if $\mathtt{e} \notin \mathtt{RSR}(\mathcal{R}\langle\Gamma,\Delta\rangle)$, then $\mathtt{e} \notin \mathtt{RSR}(\mathcal{R}\langle\Theta_1,\Lambda_1\rangle)$ or ... or $\mathtt{e} \notin \mathtt{RSR}(\mathcal{R}\langle\Theta_n,\Lambda_n\rangle)$. By Theorem 135 above, it follows that the metainference from $\{\Theta_1 \succ \Lambda_1,\ldots,\Theta_n \succ \Lambda_n\}$ to $\{\Gamma \succ \Delta\}$ is locally valid in the strong Kleene formulation of ST. ■

Proposition 139. *Suppose that the only constraint on implication-space models is that they be conic, and let* $[\![\Gamma]\!] = \mathtt{G}$ *and* $\mathtt{G}^+ = \{\mathtt{g}^+ \mid \mathtt{g} \in \mathtt{G}\}$ *and* $\mathtt{G}^- = \{\mathtt{g}^- \mid \mathtt{g} \in \mathtt{G}\}$ *and analogously for* $[\![\Delta]\!] = \mathtt{D}$. *Then:*

$\Gamma \models_{LP} \Delta$ *if and only if* $(\bigwedge \mathtt{G})^- \preceq \mathtt{D}^-$.
$\Gamma \models_{K3} \Delta$ *if and only if* $(\bigvee \mathtt{D})^+ \preceq \mathtt{G}^+$.
$\Gamma \models_{ST} \Delta$ *if and only if* $\star \preceq \mathtt{G}^+, \mathtt{D}^-$.
$\Gamma \models_{TS} \Delta$ *if and only if* $(\bigwedge \mathtt{G})^-, (\bigvee \mathtt{D})^+ \preceq \mathtt{e}$.

Proof. This follows from Fact 137 and Lemma 138. ■

We now turn to the relation between implication-space semantics and work related to analytic metaphysics, in particular Correia's logic of factual equivalence and the notions of essence, grounding, and generalized identity.

Definition 140 (Essential truth-maker models)**.** An essential truth-maker model of a truth-maker model $\mathcal{M}$ is a truth-maker model $\mathcal{M}'$ that is like $\mathcal{M}$ in its consequence relations but in which every truth-maker of A is a truth-maker of B if, for all $\langle\Gamma,\Delta\rangle$, if $\Gamma, A \models_{TM} \Delta$, then $\Gamma, B \models_{TM} \Delta$, and every falsity-maker of A is a falsity-maker of B if, for all $\langle\Gamma,\Delta\rangle$, if $\Gamma \models_{TM} B, \Delta$, then $\Gamma \models_{TM} A, \Delta$.

Essential truth-maker models, in effect, are the result of ignoring all the differences between states in a truth-maker model that do not matter for the possibility or impossibility of states.

Lemma 141. *There is a truth-maker model such that* $\models_{TM} = X$ *in that model if and only if there is an essential truth-maker model such that* $\models_{TM} = X$ *in that essential model.*

Proof. The right-to-left direction is immediate. For the left-to-right direction, take a truth-maker model $\mathcal{M}$ and add the following stipulation to it: for any state w that is a fusion of truth-makers for each sentence in a set Γ and of falsity-makers for each sentence in a set Δ, for all sets of states T and S, if (if $\forall t \in T(t \Cup w \notin S^\diamond)$, then $\forall s \in S(s \Cup w \notin S^\diamond)$), then $S \subseteq T$. This obviously does not change the consequence relation of $\mathcal{M}$. And the resulting model is an essential model of $\mathcal{M}$. To see this, suppose

that, for all $\langle \Gamma, \Delta \rangle$, if $\Gamma, A \models_{TM} \Delta$, then $\Gamma, B \models_{TM} \Delta$. Take any state w and let $\langle \Gamma, \Delta \rangle$ be such that w is a fusion of truth-makers for every sentence in Γ and falsity-makers for every sentence in Δ. Since if $\Gamma, A \models_{TM} \Delta$, then $\Gamma, B \models_{TM} \Delta$, it follows that if $\forall t \in |A|^+ (t \Cup w \notin S^\diamond)$, then $\forall s \in |B|^+ (s \Cup w \notin S^\diamond)$. And since w was arbitrary our stipulation implies that $|B|^+ \subseteq |A|^+$. For the condition regarding falsity-makers, suppose that, for all $\langle \Gamma, \Delta \rangle$, if $\Gamma \models_{TM} B, \Delta$, then $\Gamma \models_{TM} A, \Delta$. Then, for any w, if $\forall t \in |A|^+ (t \Cup w \notin S^\diamond)$, then $\forall s \in |B|^+ (s \Cup w \notin S^\diamond)$. So $|B|^- \subseteq |A|^-$. ■

This means that the nonessential truth-maker models are, in one sense, superfluous: Truth-maker theory can codify a consequence relation without the nonessential models just in case it can codify the consequence relation with the nonessential models. That is why we call these truth-maker models nonessential and the others essential truth-maker models.

Lemma 142. *Let* $[\![A]\!] = \mathtt{a}$ *and* $[\![B]\!] = \mathtt{b}$. *We have* $\mathtt{a}^+ \preceq \mathtt{b}^+$ *in all implication-space models in a set m, if and only if, in all the essential truth-maker models of truth-maker models whose consequence relation coincides with a consequence relation of a model in m, every truth-maker of A is a truth-maker of B. And* $\mathtt{a}^- \preceq \mathtt{b}^-$ *in all implication-space models in a set m, if and only if every falsity-maker of A is a falsity-maker of B in the corresponding truth-maker models.*

Proof. Left-to-right: We know from Theorem 83 that for every implication-space model, there is a truth-maker model with the same consequence relation, and vice versa. Suppose that $\mathtt{a}^+ \preceq \mathtt{b}^+$ in all implication-space models in m, that is, $\mathtt{RSR}(\mathtt{a}^+) \subseteq \mathtt{RSR}(\mathtt{b}^+)$. So, in all these models, for all $\langle \Gamma, \Delta \rangle$, if $\Gamma, A \mathrel{|\!\sim} \Delta$, then $\Gamma, B \mathrel{|\!\sim} \Delta$. Hence, in all the corresponding truth-maker models, for all $\langle \Gamma, \Delta \rangle$, if $\Gamma, A \models_{TM} \Delta$, then $\Gamma, B \models_{TM} \Delta$. In an essential truth-maker model, this ensures that every truth-maker of A is a truth-maker of B. Suppose that $\mathtt{a}^- \preceq \mathtt{b}^-$ in all implication-space models in m, that is, $\mathtt{RSR}(\mathtt{a}^-) \subseteq \mathtt{RSR}(\mathtt{b}^-)$. So, in all these models, for all $\langle \Gamma, \Delta \rangle$, if $\Gamma \mathrel{|\!\sim} A, \Delta$, then $\Gamma \mathrel{|\!\sim} B, \Delta$. Hence, in all the corresponding truth-maker models, for all $\langle \Gamma, \Delta \rangle$, if $\Gamma \models_{TM} A, \Delta$, then $\Gamma \models_{TM} B, \Delta$. In an essential truth-maker model, this ensures that every falsity-maker of A is a falsity-maker of B in the corresponding truth-maker models.

Suppose that in all the essential truth-maker models for the models m, every truth-maker of A is a truth-maker of B. Hence, in all these models, for all $\langle \Gamma, \Delta \rangle$, if $\Gamma, A \models_{TM} \Delta$, then $\Gamma, B \models_{TM} \Delta$. So, in all the implication-space models, for all $\langle \Gamma, \Delta \rangle$, if $\Gamma, A \mathrel{|\!\sim} \Delta$, then $\Gamma, B \mathrel{|\!\sim} \Delta$. Therefore, $\mathtt{a}^+ \preceq \mathtt{b}^+$. And analogous reasoning shows that if every falsity-maker of A is a falsity-maker of B in the corresponding essential truth-maker models, then $\mathtt{a}^- \preceq \mathtt{b}^-$ in all implication-space models in a set m. ■

Proposition 143. *If the only constraint on models is that they obey the semantic clauses for the implication-space semantic of* NMMS *(Definition 70), then Correia's logic and his dual system are related to implicational roles as follows. For all* $[\![A]\!] = \mathtt{a}$ *and* $[\![B]\!] = \mathtt{b}$*:*

$\vdash_c A \approx B$ *if and only if, in all models,* $\mathtt{a}^+ \preceq\succeq \mathtt{b}^+$*, and so* $\mathtt{a}^+ = \mathtt{b}^+$*.*

$\vdash_{dc} A \approx B$ *if and only if, in all models,* $\mathtt{a}^- \preceq\succeq \mathtt{b}^-$*, and so* $\mathtt{a}^- = \mathtt{b}^-$*.*

Proof. Correia uses, in effect, the same truth-maker models that we use for NMMS.[26] Moreover, Correia (2016) has shown that $\vdash_c A \approx B$ holds just in case A and B have the same truth-makers in all truth-maker models (and so merely in virtue of logical form), and it is easy to see that $\vdash_{dc} A \approx B$ holds just in case A and B have the same falsity-makers in all truth-maker models (see Hlobil, 2022a). By Lemma 141, something holds in all truth-maker models iff it holds in all essential truth-maker models. And by Lemma 142, A and B have the same truth-makers in all essential truth-maker models iff $\mathtt{a}^+ \preceq\succeq \mathtt{b}^+$. And A and B have the same falsity-makers in all essential truth-maker models iff $\mathtt{a}^- \preceq\succeq \mathtt{b}^-$. ■

Fact 144. *Factual Essence: Its being the case that A is what it is for it to be the case that B in full, in virtue of truth-functional logical form (Correia and Skiles, 2019, 651), if and only if, for* $[\![A]\!] = \mathtt{a}$ *and* $[\![B]\!] = \mathtt{b}$*, we have* $\mathtt{a}^+ = \mathtt{b}^+$ *in all implication-space models. And its being the case that A is in part what it is for it to be the case that B, in virtue of truth-functional logical form (Correia and Skiles, 2019, 651), if and only if , for* $[\![A]\!] = \mathtt{a}$ *and* $[\![B]\!] = \mathtt{b}$*, there is some* $\mathtt{c}^+$ *such that* $\mathtt{a}^+ \sqcup \mathtt{c}^+ = \mathtt{b}^+$ *in all implication-space models.*

Proof. Correia and Skiles (2019, 651) use Correia's logic to suggest that: its being the case that A is what it is for it to be the case that B in full, in virtue of truth-functional logical form, if and only if, $\vdash_c A \approx B$ in virtue of truth-functional logical form. We know from Proposition 143 that $\vdash_c A \approx B$ holds in virtue of truth-functional logical form just in case, for $[\![A]\!] = \mathtt{a}$ and $[\![B]\!] = \mathtt{b}$, we have $\mathtt{a}^+ = \mathtt{b}^+$ in all implication-space models.

Correia and Skiles (2019, 651) also use Correia's logic to suggest that: Its being the case that A is in part what it is for it to be the case that B iff: there is some C such that its being the case that B is for it to be the case that both A and C. Using Correia's logic, this holds in virtue of truth-functional logical form just in case $\vdash_c B \approx A \wedge C$. It follows from Proposition 143 that $\vdash_c B \approx A \wedge C$ holds just in case, for $[\![A]\!] = \mathtt{a}$ and $[\![B]\!] = \mathtt{b}$ and $[\![C]\!] = \mathtt{c}$, we have $\mathtt{a}^+ \sqcup \mathtt{c}^+ = \mathtt{b}^+$ in all implication-space models. So, its being the case that A is in part what it is for it to be the case that B holds, in virtue of truth-functional logical form, iff there is some $\mathtt{c}^+$ such that $\mathtt{a}^+ \sqcup \mathtt{c}^+ = \mathtt{b}^+$ in all implication-space models. ■

Fact 145. *Strict Full Grounding: Its being the case that $A_1, ..., A_n$ makes it the case that B, in virtue of truth-functional logical form (Correia and Skiles, 2019, 655), if and only if, in all implication-space models, for $[\![A_i]\!] = \mathtt{a}_i$ and $[\![B]\!] = \mathtt{b}$, (i) for some $[\![C]\!] = \mathtt{c}$, we have $(\mathtt{a}_1^+ \sqcup ... \sqcup \mathtt{a}_n^+) \sqcap \mathtt{c}^+ \sqcap (\mathtt{a}_1^+ \sqcup ... \sqcup \mathtt{a}_n^+ \sqcup \mathtt{c}^+) \preceq\succeq \mathtt{b}^+$ and (ii) $\forall 1 \leq i \leq n$ there are no $[\![D]\!] = \mathtt{d}$ and $[\![E]\!] = \mathtt{e}$ such that $(\mathtt{b}^+ \sqcup \mathtt{d}^+) \sqcap \mathtt{e}^+ \sqcap (\mathtt{b}^+ \sqcup \mathtt{d}^+ \sqcup \mathtt{e}^+) \preceq\succeq \mathtt{a}_i^+$.*

Proof. Correia and Skiles (2019, 655) suggest, using Correia's logic, that its being the case that $A_1, ..., A_n$ makes it the case that B (in the sense of strict full grounding), in virtue of truth-functional logical form, iff (i) for some C, we have $\vdash_c (A_1 \wedge ... \wedge A_n) \vee C \approx B$ and (ii) neither is there a D such that $\exists E(\vdash_c (B \wedge D) \vee E \approx A_1)$ nor ... nor is there a D such that $\exists E(\vdash_c (B \wedge D) \vee E \approx A_n)$. It follows from Proposition 143 that condition (i) holds iff for some C, with $[\![C]\!] = \mathtt{c}$, we have$(\mathtt{a}_1^+ \sqcup ... \sqcup \mathtt{a}_n^+) \sqcap \mathtt{c}^+ \sqcap (\mathtt{a}_1^+ \sqcup ... \sqcup \mathtt{a}_n^+ \sqcup \mathtt{c}^+) \preceq\succeq \mathtt{b}^+$ in all implication-space models. And condition (ii) holds iff $\forall 1 \leq i \leq n$ there are no $[\![D]\!] = \mathtt{d}$ and $[\![E]\!] = \mathtt{e}$ such that $(\mathtt{b}^+ \sqcup \mathtt{d}^+) \sqcap \mathtt{e}^+ \sqcap (\mathtt{b}^+ \sqcup \mathtt{d}^+ \sqcup \mathtt{e}^+) \preceq\succeq \mathtt{a}_i^+$. ■

Fact 146. *A sentence of the form "For it to be the case that A is for it to be the case that B" is true in virtue of logical form, in the sense of Elgin (2021, sec 4), if and only if, for $[\![A]\!] = \mathtt{a}$ and $[\![B]\!] = \mathtt{b}$, we have $\mathtt{a} = \mathtt{b}$, and so $\mathtt{a} \ni\Subset \mathtt{b}$, in all implication-space models.*

Proof. Elgin holds that two sentences are exactly equivalent iff they have the same truth-makers and the same falsity-makers. He illustrated his suggestion for generalized identity thus: "'To be a person is to be bound by the categorical imperative' holds just in case the verifiers and falsifiers of the claim that someone is a person are identical to the verifiers and falsifiers of the claim that she is bound by the categorical imperative" (Elgin, 2021, 9). If we think of sentences as 0-place predicates, this yields the claim that a sentence of the form "For it to be the case that A is for it to be the case that B" holds iff A and B have the same truth-makers and the same falsity-makers. And it holds in virtue of logical form if it holds in all truth-maker models. Now, if A and B have the same truth-makers and the same falsity-makers in all truth-maker models, then this holds in all essential truth-maker models. Hence, $\mathtt{a} = \mathtt{b}$, and so $\mathtt{a} \ni\Subset \mathtt{b}$, in all implication space models, for $[\![A]\!] = \mathtt{a}$ and $[\![B]\!] = \mathtt{b}$. ■

The facts stated in the section on nonlogical role inclusions are all immediate implications of the correspondences between sequent rules and implication-space semantics, which we highlight in various places in the chapter, especially in the subsection on the implication-space semantics of NMMS and in the subsection on linear logic.

Notes

1 Implication-space semantics is closely related to Girard's (1987) phase-space semantics for linear logic. The formalism that we present in this chapter is inspired by and similar to the one developed in Daniel Kaplan's (2022) PhD dissertation. However, our definitions of implicational roles and adjunction are different from Kaplan's definitions. As a result our semantic clauses for the logical vocabulary are also different. Unlike Kaplan, we define implicational roles by abstraction, using an equivalence relation among candidate implications.

2 Note that we do not ignore incompatibilities here. Rather, as in previous chapters, we encode implications and incompatibilities in a single reason relation, namely by using the empty set as a conclusion set in a good implication to encode the mutual incompatibility of the premises. For ease of exposition we sometimes call the combined relation an implication relation; but it should be kept in mind that this relation also includes all information about material incompatibilities.

3 Recall from Chapter Three that not just good implications but also mere candidate implications have ranges of subjunctive robustness, that is, sets of additions that yield good implications. Hence, $\langle\{\phi\},\emptyset\rangle$ need not be a good implication to have a range of subjunctive robustness.

4 We will allow ourselves to write $\mathtt{RSR}\,\langle X,Y\rangle$ for $\mathtt{RSR}(\langle X,Y\rangle)$, wherever this seems appropriate to avoid clutter.

5 We will allow ourselves to write $\mathcal{R}\,\langle X,Y\rangle$ for $\mathcal{R}(\langle X,Y\rangle)$, wherever this seems appropriate to avoid clutter.

6 The case of the role of an arbitrary candidate implication $\langle\Theta,\Lambda\rangle$ is analogous. Again, the role of a candidate implication is the set of sets of candidate implications that have the same range of subjunctive robustness: the set of sets of candidate implications such that, say, $\{\langle\Gamma,\Delta\rangle\}$ is in that set if and only if parallel additions of premises or conclusions to $\langle\Gamma,\Delta\rangle$ and to $\langle\Theta,\Lambda\rangle$ always either both yield good implications or both yield implications that are not good. Again, the idea is that an implicational role is an equivalence class, where the relevant equivalence relation is the relation of having the same range of subjunctive robustness. And we can also express this by saying that if $\langle\Gamma,\Delta\rangle$ and $\langle\Theta,\Lambda\rangle$ have the same implicational role, then they can be substituted for each other *salva consequentia*.

7 Cantor's Theorem ensures that, in every implication space, the cardinality of the bearers is always smaller than the cardinality of the potential conceptual contents in that implication space. After all, every set of candidate implications can serve as a range of subjunctive robustness that can define a premisory or conclusory role. So $|\mathcal{P}(\mathtt{S})|$, which is $|\mathcal{P}(\mathcal{P}(\mathtt{B})\times\mathcal{P}(\mathtt{B}))|$, is an upper bound on the number of premisory and conclusory roles. However, Cantor's Theorem immediately implies that $|\mathtt{B}| < |\mathcal{P}(\mathtt{B})\times\mathcal{P}(\mathtt{B})| < |\mathcal{P}(\mathtt{S})|$. Thus, the upper bounds on the number of conceptual contents defined by an implication frame is larger than the number of bearers in the implication frame. This holds only for the potential number of conceptual contents given by a set of bearers because there can be implication frames that do not define as many conceptual contents

as they have bearers. The trivial consequence relation over a collection of bearers, for example, defines only one single conceptual content.

8 If we wanted to think about conceptual contents as independent from particular implication frames, we could take another step of abstraction and say that conceptual contents defined by different implication spaces are equivalent if there is an entailment preserving bijection between the contents of the implication frames, perhaps as follows: Given two implication frames $\langle \mathtt{B}_1, \mathbb{I}_1 \rangle$ and $\langle \mathtt{B}_2, \mathbb{I}_2 \rangle$ and their contents $\mathbb{C}_1$ and $\mathbb{C}_2$, then the contents $\mathtt{a} \in \mathbb{C}_1$ and $\mathtt{b} \in \mathbb{C}_2$ are equivalent contents if and only if, there is a bijection, f, between subsets of $\mathbb{C}_1$ and $\mathbb{C}_2$ such that $\mathtt{P} \Vdash\!\sim \mathtt{C}$ holds in $\langle \mathtt{B}_1, \mathbb{I}_1 \rangle$ just in case $f(\mathtt{P}) \Vdash\!\sim f(\mathtt{C})$ holds in $\langle \mathtt{B}_2, \mathbb{I}_2 \rangle$ and $f(\mathtt{a}) = \mathtt{b}$.

With an equivalence relation among contents of different implication frames in hand, we could now define a notion of "absolute contents" as equivalence classes of contents with respect to this equivalence relation. However, we do not pursue this line of thought any further here. We do not pursue this thought because several issues arise at this point that we wish to set aside. For example, if there is more than one entailment preserving bijection between contents, then there are contents that are equivalent to two distinct contents. This happens if there are permutations of contents that preserve entailment. That might seem problematic. The issue is related to the question whether there are always several equally correct translations between any two languages (as Quine thought) and how we may be able to fix representation relations up to uniqueness, which we briefly touched upon in the previous chapter. In particular, when we consider bearers of implicational roles that are sentences of a language or vehicles of thought, the idea of Covariant Tracking might again be useful here. We may require, for example, that the bijection continues to hold under parallel additions of content bearers to both implication frames and that it continues to hold after parallel rational adjustments of the language in light of new evidence. Moreover, one might worry that our definition of equivalence of contents is not only too weak but also too strong because one might want to say that, for example, sentences in two languages have the same content although one of the two languages can express contents that the other language cannot express. Hence, one might choose to require a bijection between the contents of one implication frame and a subset of the contents of another implication frame. However, if the contents of the first implication frame are few and simple, then this relation might hold in cases in which we would intuitively deny that the contents are equivalent. Hence, the issue is a delicate one. We here set such difficult issues to one side again, as these issues do not seem specific to our account. Corresponding issues seem to arise in different ways for many theories of conceptual content and reason relations. And any tools that other theories might use to address them seem, in principle, available in our setting.

9 While languages have countably many sentences, there might be uncountably many worldly propositions. This is one respect in which the correspondence between truth-maker theory and implication-space semantics is even stronger than the correspondence between NMMS and implication-space semantics. However, we won't pursue such cardinality related issues any further here.

10 One might wonder how it can be that $[\![\phi]\!]^{\mathcal{M}} = |\phi|^{\mathcal{M}'}$ given that these objects can seem very different. Notice, however, that $[\![\phi]\!]^{\mathcal{M}} = \mathcal{R}(\phi) = \langle \mathcal{R}\langle\{\phi\},\emptyset\rangle, \mathcal{R}\langle\emptyset,\{\phi\}\rangle\rangle$, which is a pair of sets. And $|\phi|^{\mathcal{M}'} = \left\langle |\phi|^{+}, |\phi|^{-}\right\rangle$, which is also a pair of sets. So what the identity mentioned in this definition requires is that we let $\langle \mathcal{R}\langle\{\phi\},\emptyset\rangle, \mathcal{R}\langle\emptyset,\{\phi\}\rangle\rangle$ be $\left\langle |\phi|^{+}, |\phi|^{-}\right\rangle$.

11 Of course, the sentences interpreted by truth-maker models are also spaces of bearers of implicational roles. But it is more illuminating here to focus on worldly propositions.

12 The difference between contractive and non-contractive reason relations corresponds to the difference between monoids whose operation is idempotent and those whose operation is not idempotent.

13 Thinking in proof-theoretic terms, relaxing this constraint leads to so-called consecution calculi, in which we can encounter, for instance, tree structures of premises and conclusions. Such calculi are useful, for instance, in relevance logics and modal logics; but we will continue to ignore them here.

14 Daniel Kaplan (2022), who constructed the original implication-space semantics, in fact developed it by generalizing Girard's phase-space semantics because he was interested in reason relations that are nontransitive, nonmonotonic, and noncontractive.

15 In this and the next two paragraphs, we ignore some complications about the distinction between bearers and their roles. Taking them into account would make the explanation unnecessarily complicated and merely obscure what matters here.

16 To see why context-mixing Cut—given Containment—would preclude this, suppose that, for any X and Y, if $X, A \mathrel{|\!\sim} Y$, then $X, B \mathrel{|\!\sim} Y$. By Containment, $A \mathrel{|\!\sim} A$ and, hence, $B \mathrel{|\!\sim} A$. Now suppose that $X \mathrel{|\!\sim} B, Y$. By context-mixing Cut, it follows from the last two sequents that $X \mathrel{|\!\sim} A, Y$. So, if context-mixing Cut and Containment hold, then substitutability as premises implies the converse substitutability as conclusions.

17 Daniel Kaplan has first suggested such a notion in our context, in a series of working-papers presented to the ROLE group. He connected it to three-valued logics, but he did not develop the idea in the context of implication-space semantics, as we do here.

18 The principles below concern how structural principles for the underlying implications show up in role inclusion relations. Regarding structural principles for role inclusion itself, Reflexivity holds because one can replace any role for itself.

An analogue of the rule of Weakening for the left side holds for role inclusion: If $\mathcal{R}_1, ..., \mathcal{R}_n \preceq \mathcal{R}_k, ..., \mathcal{R}_m$, then $\mathcal{R}_0, \mathcal{R}_1, ..., \mathcal{R}_n \preceq \mathcal{R}_k, ..., \mathcal{R}_m$. This holds because $\mathtt{RSR}(\bigsqcap_{i=1}^{n} \mathcal{R}_i \sqcap \mathcal{R}_0) \subseteq \mathtt{RSR}(\bigsqcap_{i=1}^{n} \mathcal{R}_i)$, which is true because $\mathtt{RSR}(\bigsqcap_{i=1}^{n} \mathcal{R}_i \sqcap \mathcal{R}_0)$ is $\mathtt{RSR}(\bigsqcap_{i=1}^{n} \mathcal{R}_i) \cap \mathtt{RSR}(\mathcal{R}_0)$, which is obviously a subset of $\mathtt{RSR}(\bigsqcap_{i=1}^{n} \mathcal{R}_i)$. An analogue of the Weakening rule on the right side holds if the underlying implication relation is monotonic. That is, if the implicational

relation is monotonic, then if $\mathcal{R}_1, ..., \mathcal{R}_n \preceq \mathcal{R}_k, ..., \mathcal{R}_m$, then $\mathcal{R}_1, ..., \mathcal{R}_n \preceq \mathcal{R}_k, ..., \mathcal{R}_m, \mathcal{R}_{m+1}$.

The following analogue of Cut holds for implicational role inclusion: if $\mathcal{R}_1, ..., \mathcal{R}_n \preceq \mathcal{R}_0$ and $\mathcal{R}_0, \mathcal{R}_1, ..., \mathcal{R}_n \preceq \mathcal{R}_k, ..., \mathcal{R}_m$, then $\mathcal{R}_1, ..., \mathcal{R}_n \preceq \mathcal{R}_k, ..., \mathcal{R}_m$. For, the antecedent means that $\mathtt{RSR}(\bigsqcap_{i=1}^{n} \mathcal{R}_i) \subseteq \mathtt{RSR}(\mathcal{R}_0)$ and $\mathtt{RSR}(\mathcal{R}_0 \sqcap \bigsqcap_{i=1}^{n} \mathcal{R}_i) \subseteq \mathtt{RSR}(\bigsqcup_{j=k}^{m} \mathcal{R}_j)$. But, $\mathtt{RSR}(\mathcal{R}_0 \sqcap \bigsqcap_{i=1}^{n} \mathcal{R}_i) = \mathtt{RSR}(\mathcal{R}_0) \cap \mathtt{RSR}(\bigsqcap_{i=1}^{n} \mathcal{R}_i)$. Hence, $\mathtt{RSR}(\mathcal{R}_0) \cap \mathtt{RSR}(\bigsqcap_{i=1}^{n} \mathcal{R}_i) = \mathtt{RSR}(\bigsqcap_{i=1}^{n} \mathcal{R}_i)$. So $\mathtt{RSR}(\mathcal{R}_0 \sqcap \bigsqcap_{i=1}^{n} \mathcal{R}_i) = \mathtt{RSR}(\bigsqcap_{i=1}^{n} \mathcal{R}_i)$ and, thus, $\mathtt{RSR}(\bigsqcap_{i=1}^{n} \mathcal{R}_i) \subseteq \mathtt{RSR}(\bigsqcup_{j=k}^{m} \mathcal{R}_j)$. Therefore, $\mathcal{R}_1, ..., \mathcal{R}_n \preceq \mathcal{R}_k, ..., \mathcal{R}_m$.

19 From now on we use these semantic clauses; we assume Contraction by working with sets of bearers, and use set theoretic notions in their usual way again (for sets and not for multi-sets).

20 If we allowed ourselves to use open sentences and call their contents "concepts," we could, at this point, note that, plausibly, the concept of "_ is a dog" includes the concept of "_ is a mammal," thus vindicating the old idea that inclusions or containment relations among concepts often correspond to the inverse inclusions among the extensions of these concepts. Although working out this tantalizing idea goes beyond the scope of this book, having such an idea in mind can be helpful below when we turn to logics that have been suggested as logics for content.

21 These ascriptions would have to be qualified and explained in many respects to do justice to the positions of Priest or Kripke. We are here merely interested in the most straightforward and simplistic interpretation, as this works well enough for our current purposes.

22 This is related to work by Fitting (2021) on strict-tolerant logic and bilattices. One key difference, however, is that we do not use a lattice ordering to define consequence in implication-space models.

23 The connections that we point out in this subsection are closely related to the connections between these logics and admissible sequent rules (see Hlobil, 2022a).

24 Elgin's system is formulated in terms of predicates, but we can treat sentences as 0-place predicates. So this difference doesn't matter for our purposes. If we wanted the generality of Elgin's theory in our setting, we would need to include sub-sentential structure.

25 We formulate everything as double-line rules, which cuts down on clutter. What is new, relative to Tanter, are the rules G/FR, RFji, and LFji for the Species-Genus relation. These are variations on Tanter's CR and PR rules. We reject Tanter's CR-rule because in the absence of right-weakening the stipulation that if it follows that something belongs to a particular species, then it follows that it belongs to the genus is not equivalent with the stipulation that if it follows that something belongs to some species or other, then it follows that it belongs to the genus. We want to codify the latter and not the former. We reject Tanter's PR-rule because it is a formulation of the thought that something

belongs to a species if and only if it belongs to the genus and to none of the other species of that genus, which we can better express by saying that $F_j\nu$ should behave inferentially like $G\nu \wedge \neg F_1\nu \wedge \ldots \wedge \neg F_i\nu$, which is the same as $G\nu \wedge \neg(F_1\nu \vee \ldots \vee F_i\nu)$, which is equivalent to $\neg(\neg G\nu \vee (F_1\nu \vee \ldots \vee F_i\nu))$, which is equivalent to $\neg(G\nu \rightarrow (F_1\nu \vee \ldots \vee F_i\nu))$. And it is the latter version that is most easily seen in the sequent rules.

26 In particular, Correia uses the inclusive clauses and does not require bilateral propositions (our worldly propositions) to be convex, as Fine sometimes does.

6 Conclusion

Ulf Hlobil and Robert B. Brandom

In this book, we have developed three overarching and interrelated ideas. The first of these ideas is that propositional conceptual contents are articulated by open reason relations, that is, by reason relations that need be neither monotonic nor transitive. These reason relations are relations of implication and incompatibility, and they constrain the norms that govern practices of giving and asking for reasons. In virtue of standing to one another in reason relations of implication and incompatibility, things are bearers of conceptual contents—specifically *propositional* contents. Those contents can be thought of as the functional roles the sentences play in constellations of implications and incompatibilities. Declarative sentences express what can both serve as and stand in need of reasons (for and against, positive and negative). This is a version of semantic inferentialism. Our semantic inferentialism includes a functionalism about conceptual contents: propositional conceptual contents are individuated by the role they play in reason relations. We first formulated semantic inferentialism in a pragmatic metavocabulary in Chapter One.

The second idea is that we can make explicit these reason relations in terms of different metavocabularies. In particular, we can give accounts of reason relations in a pragmatic-normative metavocabulary and also in a semantic-representationalist metavocabulary. We call these kinds of metavocabulary "extrinsic" because they appeal to conceptual resources that are not provided by reason relations themselves. There are also two kinds of intrinsic metavocabularies, which use only conceptual resources provided by reason relations themselves. In particular, we can make reason relations explicit within the object language by using logical vocabulary, and we can also make them explicit in the metavocabulary of implication-space semantics. The idea is that these four metavocabularies afford us four perspectives on a common topic, namely reason relations. We thus undertook to understand reason relations as what shows up from the

perspective of all four metavocabularies. We called the roles played in reason relations that show up from all four perspectives "rational forms."

Our third overarching idea is that logic—appropriately understood—can serve as a guide to how the different perspectives on reason relations hang together. While it is a mistake to think that all reason relations are logical relations, logic is nevertheless intimately connected to reason relations. In particular, it is the essential and demarcating function of logical vocabulary to make explicit reason relations, and to do so within the language that is governed by these reason relations. This is the central thesis of logical expressivism, which we introduced in Chapter Two. More precisely, logical expressivism says that logical vocabulary is, firstly, vocabulary such that, for any language to which logical vocabulary is added, the reason relations among sentences in which logical vocabulary occurs can be computed from the reason relations among sentences in which logical vocabulary does not occur. And, secondly, logical vocabulary is necessary and sufficient to make explicit any reason relations among sentences of any language into which it is introduced. Thus, logic is universal in the sense that it can be explained independently of any particular reason relations of some non-logical base vocabulary. This allows us to investigate logical vocabulary from different perspectives without making assumptions about particular reason relations. Because of this expressive universality, logical vocabulary can serve as a focal point of our three other metavocabularies, thus making it easier for us to appreciate their connections.

In this final chapter, we look back at the results that we have reached by developing these three connected ideas. Before we do this, however, we want to repeat some general clarifications and caveats. First, our use of the term "reason" (we address doxastic, not practical reasons) differs from some current uses of the term. In particular, we think of implication as corresponding to the relation of being a reason for something, and we think of incompatibility as corresponding to the relation of being a reason against something. We do not assume, however, that if one accepts something that is a reason against something else, then one ought to reject the second thing. Nor do we assume that if one accepts a reason for something, then one may or ought to accept it. Rather, what we mean by "reason for" is that if *A* is a reason for *B*, then commitment to accept *A* precludes entitlement to reject *B*. And if *A* is a reason against *B*, then commitment to accept *A* precludes entitlement to accept *B*. Second, throughout the book we have restricted our attention to the propositional level and we have not said anything about subsentential expressions (apart from occasional asides).[1] Hence, there are many open questions about how our results generalize to subsentential expressions and, especially, quantifiers. Third, we want to stress that we think that the main force and importance of our technical results lies in the philosophically illuminating way in which they

connect different frameworks and perspectives, such as truth-maker theory and Restall's bilateralism or the generalization of Girard's phase-space semantics that is implication-space semantics. We do not understand our results as primarily contributions to mathematical propositional logic but rather as contributions to the philosophical understanding of propositional logic and how it is related to larger philosophical topics. Fourth, we have sometimes left open important issues and questions that have loomed large in philosophical logic, such as the issue of whether we should accept the structural rule of Contraction. We have, for instance, presented contractive and non-contractive versions of our sequent calculus and of implication-space semantics. Similarly, we have given short shrift to some influential problems in the theory of representation, such as the (alleged) under-determination of representation relations by functional roles. And we have not addressed many of the issues that arise in the current debate on the normativity of logic and which might be reproduced for the normativity of reason relations more generally, such as the question how (if at all) reason relations can sometimes make particular acceptances or rejections obligatory, thus not allowing subjects to remain neutral on an issue under discussion. These are all important issues that, in the fullness of time, a further development of the ideas we have presented here should address. We have been methodologically guided by the idea that we can legitimately leave open these issues, and by the hope and trust that they can ultimately be resolved in a way that is compatible with the ideas presented here.

Below we will first comment on the perspectives afforded by our four metavocabularies. Then we will turn to rational forms and how they occur in discourse, on the one hand, and in worldly states, on the other hand. Next we will rehearse our general strategies for theorizing about open reason relations in general. We will end by considering what our results mean for the understanding of logic.

6.1 A Fourfold Perspective on Reason Relations

We began this book by pointing to two traditions in the philosophy of language. One aims to understand meaning and content in terms of the norms that govern the use of language (or those that govern thought). The other aims to understand meaning and content in terms of what language (or thought) represents. We promised a reconciliation of these traditions, which are often understood as opposing each other. To explain how we have kept that promise, it is useful to think of the fourfold perspective afforded by our four kinds of rational expressive metavocabulary.

We understood the two traditions as theoretical undertakings pursued by talking about language in two different kinds of metalanguage: a pragmatic metavocabulary and a semantic metavocabulary, respectively.

They make it possible to express, to make explicit, different aspects of a single phenomenon. We can express normative relations among acts of accepting and rejecting in a pragmatic metavocabulary. And we can express the relation of worldly states making true or false what is asserted or rejected, as well as the modal relations among those states, in a semantic metavocabulary. Because the semantic metavocabulary and pragmatic metavocabulary are vocabularies for talking about the reason relations among bearers of conceptual contents, they are *meta*vocabularies. Because, and in the sense that, it is in particular the reason relations of the vocabularies for which they are metavocabularies that they address, they deserve to be called *rational* metavocabularies.

We presented an account of reason relations in a pragmatic metavocabulary in Chapters One and Two. We then showed, in Chapter Four, that this account is isomorphic to an account formulated in the semantic metavocabulary of truth-maker theory, even if we allow for failures of transitivity and monotonicity. That these two accounts share a common structure at the level of reason relations is the core of our reconciliation of the use-theoretic and the representationalist traditions in the philosophy of language.

The pragmatic-normative and the semantic-representationalist accounts codify the same reason relations, given a suitable agreement with respect to atomic sentences. Since conceptual contents are individuated by reason relations, it follows that the two accounts can capture the same conceptual contents. Both theories account for different contents by specifying exclusion relations among constellations of occurrences of these contents, and both theories posit two different ways in which contents can occur. The pragmatic-normative theory says that contents can occur in acceptances and in rejections, and that combinations of acceptances and rejections can be normatively ruled out or not. The semantic-representationalist theory says that contents can occur in worldly states that make them true and in worldly states that make them false, and that combination of truth-makers and falsity-makers can be alethically ruled out or not. We saw in Chapter Four how the explanation of the meanings of logical vocabulary that the pragmatic-normative theory gives as sequent rules and the explanation of the meanings of logical vocabulary that the semantic-representationalist theory gives as semantic clauses in truth-maker theory are isomorphic in their effects on reason relations. Indeed, it is only when one sees how the two accounts of logical vocabulary are isomorphic that the common structure of the two accounts becomes obvious.

The main differences between the theories lie in their explanations of what it means for contents to occur—in one of the two ways in which they can occur—and how and why some combinations of such occurrences are ruled out, where this ruling-out thereby exhibits different modal flavors. For instance, (the right kind of) semantic metavocabulary explains the

content expressed by non-logical base sentences by saying what features of the world those sentences represent. The sentence "The coin is made of pure copper" implies the sentence "The coin would melt at 1085 degrees Celsius" because and in the sense that it is alethically ruled out (it is impossible) that pure copper is solid at 1085 degrees Celsius, and it is incompatible with the sentence "The coin is an electrical insulator," because it is alethically ruled out for pure copper not to conduct electricity (pure copper necessarily conducts electricity). Pragmatic metavocabularies explain what is expressed by non-logical base sentences by saying what features of the discursive practices of *using* those sentences it is in virtue of which practitioners count as practically taking or treating the sentences *as* standing to one another in relations of implication and incompatibility. The sentence "The coin is made of pure copper" implies the sentence "The coin would melt at 1085 degrees Celsius" because and in the sense that it is normatively ruled out to accept the former but reject the latter (one cannot be entitled to that constellation of commitments), and it is incompatible with the sentence "The coin is an electrical insulator," because it is normatively ruled out to accept both sentences (one cannot be entitled to that constellation of commitments). The structure of these accounts is the same. What differs is the explanation of the relevant kind of ruling-out of combinations, the relevant kind of exclusion between occurrences of contents.

If we view the alethic exclusion relations between truth-makers and falsity-makers and the normative exclusion relation between acceptances and rejections as two sides of one coin, we thereby also view the pragmatic-normative and the semantic-representationalist theories of conceptual content as two sides of one coin. We can indeed view the alethic and the normative exclusion relations as two sides of one coin by holding that both pragmatic and semantic metavocabularies can be regarded as *rational* metavocabularies and, hence, as having at least one common topic, namely reason relations. Among the things they both discuss and seek to explain in their own terms are reason relations. This way to reconcile the two traditions in the philosophy of language becomes compelling if we say that what matters for having the topic of reason relations in view is the structure that is shared between the pragmatic-normative and the semantic-representationalist theories. That is metalinguistic functionalism about conceptual contents. If and insofar as this line of thought is correct, we should be interested not only in the way in which reason relations show up both in the norms that govern acceptances and rejections and in the modal relations among worldly states, but we should also be interested in the structure of reason relations in abstraction from these different ways in which they can be enmattered.

Appealing to the ways in which reason relations are enmattered gives semantic and pragmatic metavocabularies the expressive resources to provide complementary *explanations* of the reason relations of arbitrary base vocabularies. They are able to say, each in its own distinctive idiom, both what it means for some sentences to stand to others in relations of implication or incompatibility and why they do: what further facts explain them standing in just those relations. Because both of these kinds of metavocabulary appeal to conceptual resources beyond those intrinsic to the base vocabularies of which they are metavocabularies, and do so in service not just of characterizing the reason relations of those base vocabularies but of explaining them, the semantic and pragmatic metavocabularies can be called "extrinsic-explanatory" rational metavocabularies. When we consider reason relations in abstraction from the way in which they are enmattered, we move to characterizations of reason relations that do not have this explanatory dimension.

A metavocabulary for making explicit reason relations in abstraction from the different ways in which they can be enmattered is not an extrinsic-explanatory metavocabulary but an intrinsic-explicative one. This latter kind of metavocabulary for reason relations restricts itself to the conceptual resources supplied by the base vocabularies whose reason relations it characterizes. It is used to make explicit those reason relations and the conceptual contents they articulate, rather than to explain why they are as they are, or what it is for them to be as they are. The vocabulary of implication-space semantics is such an intrinsic-explicative rational metavocabulary.

The metavocabulary of implication-space semantics makes explicit reason relations in their unenmattered and pure form. In implication-space semantics, we make explicit the roles that things play in reason relations, merely in terms of those reason relations themselves. The way in which we do this is by considering the ranges of subjunctive robustness of implications, that is, the range of additions to the implication that yield or preserve a good implication. The relata of reason relations are treated as sets of implications in which only the given relatum occurs, as a premise or as a conclusion. The roles that relata of reason relations play are then defined as the equivalence classes of implications with respect to the equivalence relation of having the same range of subjunctive robustness—strictly, a pair of such roles, namely one in which the first member is the role of the implications in which only the given relatum occurs as a premise and the second is the corresponding role for the relatum occuring as a conclusion. In this way, we take reason relations and abstract from them the roles that relata play in them. We can then give an account of logical vocabulary just in terms of these roles. This account captures the common structure of the accounts of logical vocabulary given previously in semantic

and pragmatic metavocabularies, and it does so while abstracting from any particular way in which reason relations are enmattered.

Implication-space semantics provides a powerful new perspective on reason relations and conceptual contents because it allows us to move from particular bearers of roles in reason relations to those roles themselves. We can define operations on those roles, such as adjunction and symjunction, and we can define new roles in terms of old ones. These roles abstract away from any particular matter, shape, structure, or syntax that their bearers may have. For instance, two sentences that differ markedly in their syntactic structure can play the same implicational role. So can sentences in different languages, vocabularies with different sets of sentences. (In these respects, the conceptual contents of sentences as we understand them correspond to the functional classifications Wilfrid Sellars expresses with dot-quoted sentences such as •Der Tisch ist kühl• in his original inferentialist semantics.) We can interpret the pragmatic-normative and the semantic-representationalist theories in implication-space semantics, and doing so reveals that corresponding sentences and worldly propositions in the two theories can play the same implicational roles. That is, corresponding sentences and worldly propositions play the same role in their respective reason relations. They accordingly have the same conceptual content.

Implication-space semantics reveals, in this way, the common structure of our pragmatic-normative and the semantic-representationalist metavocabularies. We take this common structure to underlie the representation relation between sentences and worldly propositions. The idea is that, for instance, the sentence "It is raining," which can occur in acceptances and rejections, represents the worldly proposition of it raining, which can occur in worldly states that make "It is raining" true and in worldly states that make it false, because the sentence and the worldly proposition have the same conceptual content—that is, they have the same implicational role.

To sum up, the metavocabulary of implication-space semantics allows us to characterize reason relations in a way that brings out the common structure of explanations of reason relations in semantic and pragmatic metavocabularies. In particular, we can now understand the pragmatic-normative theory as a theory of which discursive acts users of concepts must treat as implying each other or being incompatible with each other in order thereby to treat their sentences as representing a particular worldly proposition. Specifically, the exclusion relations among discursive acts must be isomorphic to the exclusion relations among worldly states, in spite of the different kinds of modality that articulate the reason relations of the two kinds. That is, combining acceptances or rejections of the sentence with other discursive acts and combining truth-makers or falsity-makers of the worldly proposition with other states must always be such that either both combinations are ruled out or neither is. We thus have three different

perspectives on reason relations: one in terms of the norms governing thought and talk, one in terms of the possibility or impossibility of worldly states, and one that characterizes the roles that things play in reason relations abstractly. The third perspective binds the first two together, making them visible as providing different perspectives on one topic: reason relations.

Logical vocabulary provides the fourth perspective. Indeed, it has guided our path in binding the first two perspectives together in the third one. For it is the treatment of logical vocabulary that brings out in detail how the three perspectives relate to each other. Once we see how the rules for, say, the conditional or conjunction as stated in our three metavocabularies correspond to each other, the overall correspondence between the theories formulated in these metavocabularies becomes clear. This is so because logical vocabulary makes explicit the reason relations of arbitrary base vocabularies. As a result, by finding the correspondences between the theories of logical vocabulary in our three metavocabularies, we can appreciate the correspondence between the reason relations for the entire language.

Logical vocabulary is not a metavocabulary, in the sense that it does not talk *about* reason relations or its relata. Rather, it makes reason relations explicit within (an extension of) the object language itself. The vocabulary of the sequent calculus, which we use to introduce logical vocabulary as universally LX, *is* a rational metavocabulary in the strict sense. Logical vocabulary is, nevertheless, a rational metavocabulary in the sense that it allows us to form sentences whose undeniability makes explicit what implies what and what is incompatible with what. To put it differently, there are three ways in which we can commit ourselves to an implication holding between *A* and *B*. First, we can infer *B* from *A* or, equivalently, put *A* forward in response to a challenge of an acceptance of *B*. Second, we can treat the sentence "If *A*, then *B*" as undeniable. Third, we can assert: "That *A* implies that *B*" or "'*B*' follows from '*A*'." Logical vocabulary allows us to endorse reason relations in the second way. The other three kinds of metavocabulary allow us to endorse reason relations in the third way. Logical vocabulary and implication-space semantics have in common that they do not appeal to conceptual resources that are not provided by the reason relations that they make explicit. They both characterize reason relations in a way that is neutral both with respect to the nature of the exclusion among occurrences of contents that is essential to reason relations, and even to the nature of those occurrences themselves (the 'bearers' of conceptual contents). That is what makes them *intrinsic* rational metavocabularies. Semantic and pragmatic metavocabularies have in common that they have the resources to formulate substantive explanations of why "implies" or "follows from"

apply in particular cases. Because these explanations differ, it is not obvious that they are compatible, so that it can seem that theories formulated in semantic and pragmatic metavocabularies pull in different directions. Implication-space semantics binds together the perspectives provided by the semantic and the pragmatic metavocabularies, as it were, from above, while logical vocabulary binds them together from below.

This fourfold perspective on reason relations, conceptual content, and implicational roles is the core of the ideas that we have presented here. It takes some technical work to show how the accounts of logical vocabulary given in the three kinds of metavocabulary are all equivalent, but it can be done, and the work is richly rewarded. We are not merely waving our hands at the fourfold perspective just mentioned. We have shown that it can be treated with mathematical rigor. Such a rigorous treatment brings with it fruitful and specific clarifications of relations between our ideas and extant theories. Along the way, we have learned, for instance, how to add nonmonotonic implications to a sequent calculus while validating the deduction-detachment theorem and other desirable principles,[2] how to formulate strict-tolerant logic in truth-maker theory, how to formulate multiplicative-additive linear logic in terms of implicational roles, how to recover Correia's logic of factual equivalence from inclusion relations among implicational roles, and much more. We will come back to these fruits of the fourfold perspective that sits at the core of this book. Judging by these fruits, we think that our fourfold perspective has proven its worth.

Explicating, explaining, and proving the relations among the perspectives on reason relations and conceptual contents afforded by these four different kinds of metavocabulary also strengthens the roots of the tree that produces these fruits. As T.S. Eliot famously says in his *Four Quartets*, "We shall not cease from exploration, and the end of all our exploring will be to arrive where we started and know the place for the first time." We began our story by introducing the idea of *reason relations* of implication and incompatibility (via reasons for and reasons against) as one way of thinking about the fundamental metaconceptual bipolarity that also shows up both in the pragmatic opposition between acceptance and rejection and in the semantic opposition between truth and falsity. When they first came on the scene, reason relations were recommended as a way of thinking about language—a notion made more precise in the form of our metaconcept of a *vocabulary*, understood as a set of reason relations defined on a lexicon of declarative sentences. We have since seen that this metalinguistic rationalism, while productive, was too narrowly conceived. Language is not the only medium of reason relations. Not only what we say, but also what we talk *about* is essentially articulated by reason relations of modally robust inclusion and exclusion. Truth-maker semantic metavocabularies let us talk about nonlinguistic, modally robust, worldly matter-of-factual relations

of consequence and incompatibility. Specifically meta*linguistic* rationalism accordingly takes its place as one aspect of a broader meta*conceptual functionalism* about reason relations. That is the strategy of picking out and understanding reason relations functionally: by the complex role we have now seen them play as the demonstrably common topic of (the right sort of) pragmatic, semantic, logical, and implication-space conceptual-role model-theoretic metavocabularies.

6.2 Rational Forms

Looking at reason relations from the perspectives of our four rational metavocabularies reveals not only the essential structure of reason relations but also how roles in reason relations can be shared between different manifestations of those reason relations. We can think of the different manifestations of reason relations as different ways in which such roles are enmattered. We can accordingly think of roles in reason relations as forms that can inform different kinds of matter. That is the idea of rational forms, which is another thread that runs through the entire book.

The difference between the explanations of reason relations in the pragmatic metavocabulary of Chapters One and Three and in the semantic metavocabulary of Chapter Four is a difference in the kind of modality and a corresponding difference in the nature of occurrences of contents to which these two explanations appeal. The pragmatic explanation appeals to a normative exclusion relation between occurrences of contents that are discursive acts or attitudes. The semantic explanation appeals to an alethic exclusion relation between occurrences of contents that are worldly states. We suggest that we think of these differences as pertaining to the matter in which reason relations are manifested and that we think of what is shared between the two explanations as the forms that are enmattered in these different ways. The forms that are thus enmattered are what we call "*rational forms*."

Rational forms occur in discursive acts and also in worldly states. And since rational forms are nothing but the roles in reason relations that are propositional conceptual contents, it follows that conceptual contents occur in discursive acts and also in worldly states. In this sense, conceptual contents are real—not only in that conceptual contents occur in thought and talk but in the stronger sense that the structure of reality is a conceptual structure. That is, the structure of reality itself is articulated by reason relations. Our conceptual realism can aptly be denominated "bimodal" because the crucial difference between the two ways in which conceptual contents can occur is the difference between the kinds of modality characteristic of the exclusion relation between occurrences of contents. The idea is that the difference between kinds of matter in which

rational forms are enmattered comes with a difference between the kinds of modality that characterizes exclusion relations between occurrences of forms in the particular matter at hand. When rational forms are enmattered such that their occurrences are worldly states, then the relevant exclusion relation between occurrences is alethic. And when rational forms are enmattered such that their occurrences are discursive acts, then the relevant kind of exclusion relation is normative (deontic). This is bimodal conceptual realism. Because we understand material relations of consequence and incompatibility as reason relations, on our conception, to be a modal realist is to be a conceptual realist. That is, taking it that the world as it is objectively includes relations of necessary consequence and relative impossibility of the sort expressed in laws of nature and true subjunctive conditionals *is* taking it to have a *conceptual* structure in the sense we have explained and justified giving to that term.

Bimodal conceptual realism is *hylomorphic* in a recognizably Aristotelian sense. There are, however, important differences between Aristotelian notions of form and our rational forms. First, unlike Aristotelian forms, the metaphysical category to which bearers of rational forms belong is not primarily that of concrete particulars but that of proposition-like or sentence-like entities, namely entities whose occurrences stand to one another in reason relations of consequence and incompatibility. Second, something having a rational form does not constitute a concrete particular. Rather only *occurrences* of rational forms are concrete particulars, and what has a rational form can occur in two different ways. Third, rational forms are individuated by the modal relations between their occurrences and occurrences of other rational forms, and these modal relations can be of two different kinds, normative and alethic. So, where Aristotelians have, for the primary case of substances, forms whose manifestation in matter yields concrete particulars, we have bimodal forms whose manifestations can occur in two ways to yield concrete instances of proposition-like entities. What Aristotelian forms and our rational forms have in common is that the same form occurs in a thought and in what the thought represents. As already intimated in the Introduction to this book, this is broadly the same idea that Spinoza expresses in his slogan: "The order and connection of ideas is the same as the order and connection of things" (Spinoza, *Ethics* II, Prop. vii). The reason relations that define rational forms articulate both the order and connection of discursive acts (ideas) and the order and connection of worldly states (things).

At the end of Chapter Five we pointed out that the implication-space conceptual role semantics that is the native language of rational forms as such is both holistic and compositional. More complex conceptual roles are constructed from simpler ones by adjunction and symjunction. And the logic NMMS computes the reason relations of logically complex

sentences from those of logically simpler ones by (variants of) the usual connective rules of sequent calculi. The hoary and hackneyed opposition between coherence theories and correspondence theories of *truth* reflects a more fundamental distinction between holistic ("coherence") theories and representational ("correspondence") approaches to *meaning* or conceptual content. Rational forms are functionally defined to begin with by the reason relations linguistic or worldly bearers stand in to one another within their broadly subjective and objective realms, respectively. As with anything defined by its functional role in some larger relational whole, rational forms are accordingly identified and individuated holistically. But because rational forms can be shared between the two realms—the norm-governed discursive practices of language users and the subjunctively robust relations among worldly states—that holistic specification of conceptual contents builds in the dimension of correspondence that we are calling "bimodal conceptual realism." Our view is accordingly distinctive in locating the most basic correspondence between words and the world in reason relations, and thereby in rational forms or propositional conceptual contents. This fundamental representational dimension only becomes visible in terms of the relations between what is expressed in pragmatic metavocabularies specifying subjects' use of expressions in deontic normative terms and semantic metavocabularies specifying the objective world represented in alethic modal terms. It is not reducible to, or intelligible just in the terms of, the notion of representation codified in the interpretation functions that truth-maker semantic metavocabularies specify. The meta-metalinguistic character of our account of representation as the amphibiousness of rational forms between language and the world, with its essential appeal to pragmatic metavocabularies for specifying the use of linguistic expressions, operates in a metaconceptual space that, in spite of its historical credentials, is essentially richer than talk about representation or correspondence in terms of truth or meaning.

Rational forms provide a way to understand why and how what is expressed by declarative sentences is both what can be semantically evaluated as true or false and what can be pragmatically accepted or rejected. For what is expressed by declarative sentences are propositional conceptual contents or, equivalently, rational forms. If a rational form occurs positively in a worldly state, then this makes what the sentence expresses true. And if it occurs negatively, this makes what the sentence expresses false. If a rational form occurs positively in a discursive act, then this act is an acceptance of what the sentence expresses. And if the rational form occurs negatively in a discursive act, then this act is a rejection of what the sentence expresses. Thus, what is expressed by declarative sentences is what is made true or false by worldly states, and it is also what is accepted and rejected in discursive acts.

Rational forms have a dual bilaterality. They have two sides in that they can occur in two opposing ways, positively and negatively. And both of these sides can occur in two different kinds of matter, the matter of worldly states related by alethic exclusions and the matter of discursive acts related by normative exclusions. The duality of having a positive and a negative way to occur reflects the fact that rational forms are articulated by relations of reasons for and reasons against. The duality of occurring in two kinds of matter reflects the isomorphism of reason relations in pragmatic and in semantic metavocabularies. In semantic terms, we can say that the first duality corresponds to the contrast between truth and falsity, while the second duality corresponds to the contrast between representings and what is represented. The space in which rational forms can occur is the field spanned by these four corners.

The development of the notion of rational forms was the main result of Chapter Four. There we showed how we can understand the normative bilateralism of Restall and Ripley, on the one hand, and Fine's bilateralism of truth-makers and falsity-makers, on the other, as two sides of reason relations and, hence, of the rational forms they articulate. Developing this notion yielded as a by-product a truth-maker semantics for strict-tolerant logic, as well as a generalization of both strict-tolerant logic and truth-maker theory to open, nonmonotonic and nontransitive consequence relations. Along the way, we saw how proof-theoretic sequent rules correspond in a fine-grained way to the semantic clauses of truth-maker theory (and, later, to constructions in implication-space semantics). And we formulated a novel way to understand consequence in truth-maker theory in modal terms, thus forging a connection to modalism about consequence and the long tradition behind it. Developing the big philosophical idea of rational forms that are shared between representations and what is represented thus paid significant dividends at the level of technical details. The simultaneous work on the big philosophical issue and the technical details gives us a notion of rational form that has the potential to illuminate the subject-object nexus from a perspective that is not only new, but is one that we can also manipulate with mathematical precision and rigor.

6.3 Theorizing Open Reason Relations

One of our overarching aims in this book was to develop methods for theorizing open reason relations, that is, reason relations in which monotonicity and transitivity can fail. More specifically, our aim was to develop such methods in a way that is consistent with semantic inferentialism and logical expressivism. This is a thread that binds together the three ideas with which we opened this chapter.

Implication-space semantics, which we have presented in Chapter Five, is the most abstract way to theorize open reason relations. Implication-space semantics abstracts away from the duality of two kinds of matter in which rational forms can occur and characterizes rational forms independently of the matter in which they can occur. This requires abstracting away from the substantive concepts *extrinsic* to the base vocabularies for which they are metavocabularies that permit semantic and pragmatic *explanations* of reason relations. The implication-space metavocabulary provides a model-theoretic semantics for the conceptual roles bearers of conceptual contents play in virtue of standing to one another in reason relations of implication and incompatibility. It is a reason-based inferentialist semantics, rather than a truth-based representational semantics such as truth-maker semantics.

If the devil is in the details of different kinds of matter, then the angels will be found working with rational forms at the more abstract (intrinsic) level. While we characterize reason relations in the pragmatic and semantic metavocabularies in terms of a particular modal exclusion relation, implication-space semantics includes a materially undifferentiated and primitive distinction between good and bad implications. It is materially undifferentiated in the sense of being amphibious between the characterizations in a normative pragmatic and in an alethic modal semantic metavocabulary. While we characterize occurrences of contents as acceptances and rejections and as truth-makers and falsity-makers in the pragmatic and semantic metavocabularies, respectively, implication-space semantics includes a materially undifferentiated and primitive structural bipolarity of positive (premisory) and negative (conclusory) occurrences of conceptual contents. And while we characterize combinations of occurrences of contents as discursive positions or as fusions of states in the pragmatic and semantic metavocabularies, respectively, implication-space semantics includes an undifferentiated and primitive mode of combining implications. As the pure (intrinsic) metavocabulary of conceptual roles, inferentialist implication-space semantics is our candidate for the language of the angels, who immediately grasp the abstract rational forms shared by worldly states on the objective side of the intentional nexus and discursive acts on the subjective side. For us fallen creatures, the intrinsic-explicative metavocabularies of implication-space semantics and logic are tools for coming closer to angelic rational self-consciousness. For they allow us to become reflective and critical—in our limited way—with respect to reason relations, independently of the way in which they are enmattered.

The three kinds of abstraction just sketched allow implication-space semantics to capture the essential structure of reason relations. And we saw in Chapter Five that this is the structure of a commutative monoid defined on a set of pairs, together with a partition of that set of pairs.[3] The primitive distinction between good and bad implications shows up as the partition

of the monoid set by the particular subset. That conceptual contents can occur in two different ways shows up in the fact that the elements of the monoid are pairs. And the primitive way to combine implications shows up as the monoidal operation. These are the basic ingredients one needs to build an implication-space model in which one can interpret bearers of conceptual contents, such as sentences.

We can think of the construction of such a model as happening in two stages. At the first stage, what is interpreted, as well as the building blocks of the interpretants, are implications. Thus, implication-space semantics is an inferentialist semantics not only in that its semantic interpretants are implicational roles but also in that what is interpreted are, in the first instance, candidate implications. The semantic interpretation of a candidate implication is its equivalence class with respect to its range of subjunctive robustness, that is, the class of implications that have the same range of subjunctive robustness. Our monoidal structure provides the resources to define ranges of subjunctive robustness because we can say that the range of subjunctive robustness of an element of our monoid set is the set of elements such that when combined by the monoidal operation with the target element yield a result in the distinguished subset. Therefore, as long as candidate implications form a commutative monoid of pairs with a particular subset of good implications, we can define ranges of subjunctive robustness and, hence, define the semantic interpretants of implications.

The ranges of subjunctive robustness of candidate implications are their "goodness" conditions, as truth conditions are the "goodness" conditions of sentences in truth-conditional semantics. For a candidate implication to be good in the reasons-first semantic setting is for its premises to provide reasons for its conclusion, while for a sentence to be good in the truth-first semantic setting is for it to be true. The advance from a conception of semantic goodness to a conception of meaning is the advance to consideration of circumstances under which a reason relation or sentence *would* be good. In truth-maker theory, the circumstances are sets of worldly states. In implication-space semantics, the circumstances are additional premises and additional conclusions that would make or keep the implication good. The semantic interpretation of the implication is then the class of implications that have the same conditions of semantic goodness.

At the second stage, we interpret particular bearers of conceptual contents, such as sentences. The semantic interpretant of a sentence (or other particular bearer of a conceptual content) is its implicational role, and the implicational role of a sentence is a pair whose first element is the sentence's premisory role and whose second element is the sentence's conclusory role. The premisory role is the semantic interpretant of the candidate implication that is a bare occurrence of the sentence as a premise,

and the conclusory role is the semantic interpretant of the candidate implication that is a bare occurrence of the sentence as a conclusion.

In this way, implication-space semantics isolates the roles that its relata play in any given reason relation. It is a way of arriving at the implicational roles of things that stand in reason relations by means of mathematical abstraction. Crucially, this works for structurally open reason relations as well as for topologically closed ones. Indeed, it works not only for reason relations in which Monotonicity and Transitivity can fail but also for reason relations in which Contraction and Reflexivity can fail. It is accordingly an extremely versatile and flexible tool for codifying the implicational roles that the relata in open reason relations play.

These aspects of implication-space semantics apply to open reason relations even before we consider logical vocabulary. The reason relations implication space-semantics codifies need to be closed neither in the sense of being a topological closure operator nor in the sense that explicitation of implications is inconsequential. That is, these reason relations need to obey neither Monotonicity, Cut, and Reflexivity, nor the weaker principles of Cautious Monotonicity and Cumulative Transitivity. So if Γ implies A, then adding A to Γ can change what is implied in both directions: by losing consequences and by adding consequences. That is, if—in addition to A—the premises Γ also imply B but do not imply C, then the combination of Γ and A can nevertheless fail to imply B, and it can imply C. So making the implications of a set of premises explicit in the sense of adding them to the premise-set can have substantive effects on what is implied. In this sense, the kind of inference that consists in acknowledging commitment to implications of one's commitments by taking them as further explicit premises is a creative act. It may open up paths to implications one has not foreseen, and close alternative ways to develop one's commitments. Equally reasonable paths of explicitly endorsing implications of the same starting point can lead to incompatible end points. As already intimated in Chapter Two, this rational hysteresis contains the seed of the fact that it is essential to rational beings to have not merely a past but a history. There need not be any single end-point on which all fully rational creatures will converge when they draw out the consequences of a common starting point, and whether the position at which some rational creatures arrived in this way is reasonable depends in part on the path that led them there. Rationality can be a matter of how one got to where one is, not merely in the psychological sense of depending on the particular reasons for which one undertook a given commitment, but in the sense that different reasons are available on different paths from a starting point to an end point.

Developing the tools for theorizing open reason relations in a formally rigorous way has proven to be fruitful in unexpected ways. Only by doing so could we use Girard's phase-space semantics for linear logic as a stepping

stone to the essential structure of reason relations. And it is only once one sees that this structure is the structure of a monoid defined on a set of pairs with a partition on it that one can clearly see this structure in the theories formulated in pragmatic and semantic metavocabularies. As a by-product of this insight, it becomes clear that Girard's phase-space semantics uses the notion of the dual of facts to do two tasks that we ought to keep separate, namely the task of giving an account of negation and the task of defining the ranges of subjunctive robustness that define implicational roles. And we can prevent conflating these two tasks by using pairs as the elements of our monoid set, and letting the structure of pairs underlie our account of negation. It was this formal investigation that cleared the path to a deeper philosophical understanding of open reason relations.

6.4 The Logics of Open Reason Relations

One of many uses that one can make of implication-space semantics is to compute the implicational roles of arbitrary logically complex sentences from the implicational roles of logically atomic sentences of any base vocabulary—even when the reason relations of the base vocabulary, and so its (conservative) logical extension, are radically substructural. We do this by defining the operations of adjunction and symjunction on implicational roles. Adjunction combines positive or negative occurrences of conceptual contents in a single implication and returns the implicational role of that combined implication. It is a way to look at the joint effects of several occurrences of different conceptual contents in a single implication. Symjunction looks at the intersection of the ranges of subjunctive robustness of roles and returns the implicational role whose members have that intersection as their range of subjunctive robustness. It is, as it were, a way to look at the overlap of, or at what is shared between, the implicational roles of several implications. Thus, adjunction combines occurrences of conceptual contents in single implications and takes the implicational role of the result, whereas symjunction combines the roles of several implications by looking at the intersection of their ranges of subjunctive robustness. With these operations on the table, we can say, for instance, that the conclusory implicational role of a conditional is the adjunction of the premisory implicational role of the antecedent and the conclusory implicational role of the consequent. And we can say that the implicational role of a negation is the result of swapping the premisory and the conclusory implicational roles of the negatum.

There is a close connection between the way in which the implicational roles of logically complex sentences are computed in model-theoretic implication-space semantics and the proof-theoretic sequent rules for the logical connectives. The specification of the premisory role of a logically

complex sentence corresponds to the right-rule for the principal connective of the sentence, and the specification of the conclusory role corresponds to the left-rule. Adjunction corresponds to sequent rules with just one top sequent, and symjunction corresponds to rules with several top sequents. If an adjunction or symjunction operates on a premisory role, then the sentence with that role occurs on the left side in the (or a) top sequent. And if an adjunction or symjunction operates on a conclusory role, then the sentence with that role occurs on the right side in the (or a) top sequent. In this way, one can read off sequent rules from clauses for implicational roles, and one can read off clauses for implicational roles from sequent rules. Indeed, our brief discussion of multiplicative-additive linear logic in Chapter Five showed that this mapping can be extended from the Ketonen sequent rules to sequent rules for the multiplicative and additive connectives of linear logic.

There is also a close connection between the way in which the implicational roles of logically complex sentences are computed in implication-space semantics and the semantic clauses of truth-maker theory. The specification of the premisory role of a logically complex sentence corresponds to clauses for verifiers in truth-maker theory, and the specification of the conclusory role corresponds to the clauses for falsifiers. Adjunction corresponds to truth-maker clauses in which the states that verify or falsify the compound sentence are fusions of two states, and symjunction corresponds to truth-maker clauses that specify the states that verify or falsify the compound sentence by a disjunction. Appeals to verifiers on the right-hand side of truth-maker clauses correspond to appeals to premisory roles, and appeals to falsifiers on the right-hand side of truth-maker clauses correspond to appeals to conclusory roles.

There is an equally simple correspondence for the structural rules. The Weakening rule in the sequent calculus corresponds to Downward-Closure in truth-maker theory and to the requirement that all ranges of subjunctive robustness be maximal in implication-space semantics. The principle of Containment in the sequent calculus corresponds to Exclusivity in truth-maker theory and it corresponds to the principle that pairs with overlapping elements are good implications in implication-space semantics. The Cut rule in the sequent calculus corresponds to Exhaustivity in truth-maker theory, and in implication-space semantics it corresponds to the principle that if an implication's range of subjunctive robustness includes the symjunction of the premisory and the conclusory role of some one thing, then the implication is good.

By appreciating these correspondences, the underlying shared structure of reason relations becomes visible in all of the theories. Implication-space semantics captures this structure in the clearest and most abstract way, by allowing us to talk directly about roles in reason relations, rather

than merely the relata in reason relations that play these roles. However, once one sees the common structure, it is easy to identify these roles in each of the different theories. As already intimated, the roles of logically complex sentences guide the way here because they make explicit arbitrary material reason relations. Once we see the correspondences in the case of logically complex sentences, the general isomorphism between the structures becomes difficult to overlook.

In a final additional step of abstraction, we can define inclusion relations between implicational roles in terms of set-theoretic inclusion relations between their ranges of subjunctive robustness. The idea is that, relative to some condition on implication-space models, an implicational role is included in another implicational role if and only if, in all implication-space models that meet the condition, the range of subjunctive robustness of the former role is a subset of the range of subjunctive robustness of the latter role. If this holds, one can (in all these models) replace bearers of the former role by bearers of the latter role *salva consequentia*, that is, without turning a good implication into a bad one. Thus, implicational role inclusion relations codify relations of substitutability of bearers of these roles in implications *salva consequentia*.

The vocabulary of such role inclusion relations is an intrinsic rational metametavocabulary, built on top of the implication-space model-theoretic semantics for conceptual roles. It provides the expressive power to make explicit metainferential reason relations among implicational roles. In particular, we can recover the notion of local metainferential validity in terms of implicational role inclusions. This offers an illuminating new semantic perspective on the relations among a variety of well-studied logics. The paraconsistent logic of paradox (LP) shows up as the logic of conclusory role inclusions. That is, A implies B in LP if and only if the conclusory role of A is included in the conclusory role of B, in all conic implication-space models. And the dual paracomplete strong Kleene logic (K3) shows up as the logic of premisory implicational role inclusions. That is, A implies B in K3 if and only if the premisory role of A includes the premisory role of B, in all conic implication-space models. And we can formulate strict-tolerant and tolerant-strict logic in a similar way. Moreover, Correia's logic of factual equivalence results from our formulation of K3 by dropping the constraint that models be conic but require that the implication goes in both directions. And Correia's dual logic results from our formulation of LP in the analogous way. Indeed, we can even recover analogues of the notions of essence and grounding in terms of implicational role inclusions.

The laws of all these familiar logics show up in implication-space semantics as licenses to substitute bearers of certain roles *salva consequentia*. For that is what inferential role inclusions are. This is

how what is currently studied under the heading of metainferential validity emerges in implication-space semantics. However, implication-space semantics is much more flexible and powerful than the usual three-valued semantic frameworks in which many of the above-mentioned logics are usually studied. In contrast to the usual three-valued setting, we can codify not only nontransitive but also nonmonotonic consequence relations in implication-space semantics. Similarly, we can recover Correia's logics, which cannot be formulated in a standard three-valued setting. Similarly, the connection to truth-maker theory and to linear logic cannot be appreciated from the perspective of a standard three-valued setting. By contrast, if one adopts our fourfold perspective on reason relations, all these rich and interesting connections become visible. Thus, we think that we have not only found a way to theorize open reason relations but a genuinely eye-opening way to do so.

Looking back, we can see the structure revealed by implication-space semantics in the NMMS sequent calculus from Chapter Three—a structure that cannot easily be seen by looking merely at that calculus. However, the sequent calculus brings out very clearly the idea of treating the material implications of nonlogical vocabulary as the given basis to which we then add logical vocabulary. For in the sequent calculus, we treat the material implications (and incompatibilities) as axioms, which are then closed under the sequent rules. In this way, the complete logically extended vocabulary (lexicon of logically complex sentences plus reason relations among them) can then be computed from the base vocabulary. That we can close arbitrary sets of nonlogical axioms under these sequent rules is what it means that the logically extended vocabulary can be *elaborated from* any arbitrary base vocabulary.

One thing one can reasonably mean by the "logic" of this logical extended vocabulary is the implications and incompatibilities that hold in every logical extension of any base vocabulary that obeys Containment. These are sequents such as $\Gamma, A \to B, A \mid\!\sim B, \Delta$ that hold no matter what implications hold in the base vocabulary, simply in virtue of the rules for determining the reason relations that logically complex sentences stand in, together with Containment. This logic is just classical logic. So, in this sense, the logic of open reason relations that we are suggesting is classical propositional logic.

Classical logic is special because the constraint of Containment can plausibly be taken to be an essential part of the structure of conceptual contents. For, according to the view that we are suggesting, conceptual contents have essentially two opposing sides, and Containment can be understood as saying that the two sides of any genuine conceptual content always exclude each other, no matter in what context they occur. Thus, Containment can plausibly be understood as reflecting the opposition

between the two sides of contents. If one accepts this idea, then classical logic is a formulation of the reason relations that hold in all logically extended vocabularies. This is what we think the tradition got right in putting classical logic at the center of logic and the study of reason relations. It would be a mistake, however, to ignore the richer open reason relations within which classical logic emerges as a minimal part.

The sequent calculus of Chapter Three also suggests several different clear ways to understand the claim that logical vocabulary makes explicit arbitrary reason relations. One way to understand this central claim of logical expressivism is this: A sentence ϕ makes explicit that Θ is a reason for Λ if and only if, for all Γ and Δ, we have $\Gamma \mathrel{|\!\sim} \phi, \Delta$ just in case $\Gamma, \Theta \mathrel{|\!\sim} \Lambda, \Delta$. And the sentence ϕ makes explicit that Θ is a reason against Λ if and only if, for all Γ and Δ, we have $\Gamma \mathrel{|\!\sim} \phi, \Delta$ just in case $\Gamma, \Theta, \Lambda \mathrel{|\!\sim} \Delta$. We saw in Chapter Three that the following conditions jointly ensure that logical vocabulary allows us to form sentences that can make arbitrary reason relations explicit in that sense.

Deduction-Detachment (DD) Condition on Conditionals:
$\Gamma \mathrel{|\!\sim} A \to B$ if and only if $\Gamma, A \mathrel{|\!\sim} B$.

Incoherence-Incompatibility (II) Condition on Negation:
$\Gamma \mathrel{|\!\sim} \neg A$ if and only if $\Gamma, A \mathrel{|\!\sim}$ if and only if $\Gamma \# A$.

Antecedent-Adjunction (AA) Condition on Conjunctions:
$\Gamma, A, B \mathrel{|\!\sim} \Delta$ if and only if $\Gamma, A \wedge B \mathrel{|\!\sim} \Delta$.

Succedent-Summation (SS) Condition on Disjunctions:
$\Gamma \mathrel{|\!\sim} A, B, \Delta$ if and only if $\Gamma \mathrel{|\!\sim} A \vee B, \Delta$.

That these conditions hold can be appreciated from the perspective of truth-maker theory and implication-space semantics. In the sequent calculus setting, however, that these conditions hold follows immediately and transparently from the sequent rules of NMMS (and the variants we have presented), and the fact that they are invertible. Insisting that these conditions hold, even in logically extended open reason relations, yields results that are familiar in many ways but unorthodox in other ways. Some unorthodox consequences are, for instance, that premises can imply a conditional and its antecedent without implying its consequent (that is, *meta-modus-ponens* can fail), and that sentences that are equivalent in classical logic cannot always be substituted for each other *salva consequentia*. We hope to have shown over the course of the book that accepting such unorthodox consequences not only gives us logical vocabulary that can make arbitrary reason relations explicit, but also opens the door to new and illuminating perspectives on representation, the

relations between semantic and pragmatic metavocabularies, the essential structure of reason relations, and a rich network of surprising connections among familiar logical theories.

In Chapter Three, we also offered another sense in which logical vocabulary makes explicit the reason relations of arbitrary base vocabularies. The idea is that every reason relation among logically complex sentences reflects—and thereby expresses—specific reason relations among sentences of the base vocabulary. We saw that the version of NMMS in which Contraction holds fulfills this task in one direction: For any sequent of the NMMS-extension of any base vocabulary, the sequent is derivable if and only if a unique and easily specifiable set of sequents is in the base consequence relation. In order to fulfill this task also in the other direction, we must move to the version of NMMS in which Contraction fails. In that version, for every arbitrary set of sequents, X, that can be formulated in the language of the base vocabulary, there is a single sequent in the logically extended vocabulary such that this sequent is derivable if and only if all the sequents in X are in the base consequence relation. So, in the noncontractive version of NMMS every collection of base sequents can be represented by a single sequent in the logical extension. This result also applies, of course, to the implication-space formulation of the noncontractive version of NMMS. Hence, the theories that we have put forward in this book treat logical vocabulary in such a way that it also makes arbitrary reason relations explicit in the second of the two senses (though this holds only in one direction for the contractive versions).

Putting the two points together, we can say: NMMS allows us to elaborate logical vocabulary from arbitrary base vocabularies, and the logical vocabulary of NMMS makes arbitrary reason relations of any given base vocabulary explicit. We take the combination of these two facts to mean that the theories that we have presented in this book live up to the ideal of logical expressivism. They underwrite the claims that logical vocabulary is *el*aborated from and *ex*plicative of (LX for) the reason relations of arbitrary base vocabularies. Because it can extend any and every base vocabulary (that meets the most minimal conditions), the ideal logic is *universally* LX. The theories we have presented approach this ideal of being universally LX.

As a by-product of developing these ideas in a rigorous way, we learned how to introduce operators that make explicit when implications hold monotonically (in a generally nonmonotonic setting). And we also showed how the strategy can be generalized to make explicit both the classicality of an implication (in a generally nonclassical setting) and when contraction holds (in a generally noncontractive setting).

To sum up, there is a sense in which the logic of open reason relations is simply classical logic. However, our fourfold perspective on reason

relations revealed many other familiar logics in the structure of reason relations. In particular, since Cut can fail, we can think of the classical consequence relation that emerges as the logic of open reason relations as really being a formulation of strict-tolerant logic. Moreover, the logic of paradox, for instance, is the logic of conclusory role inclusions, in conic models, and Correia's dual logic results if we look at all models and not just the conic ones. And strong Kleene logic stands in the analogous relation to Correia's logic for premisory role inclusions. Furthermore, if we move from Containment to Reflexivity and allow for failures of Contraction, then the structure of implication-space semantics, in effect, coincides with the structure of phase-space semantics for multiplicative additive linear logic, and we can define the linear connectives in that setting. All of these connections hold despite the fact that we are far less exclusive in the reason relations that we consider than virtually all of these other logics. The logic of paradox and strong Kleene logic are monotonic, Correia's logics and linear logic are transitive, and none of these theories considers material implications and incompatibilities. None of these logics is *the* logic of open reason relations. Rather, they all capture specific aspects of open reason relations. Seeing open reason relations from the perspective of the four rational metavocabularies in the way we have suggested provides a clear and detailed picture of the different aspects that these logics capture.

6.5 Final Remarks

This is a book about reasons and, in particular, reason relations. It is animated by a kind of rationalism about philosophy, language, logic, and meaning that consists in the interlocking convictions that:

- Concern with the nature of reasons is coeval with philosophy, and abides at its core.
- Language properly comes to be a central focus of philosophical attention insofar as it is understood as a medium of reasons.
- The methodological promise of logic as a tool for philosophical understanding stems from the idea that it expresses the essential form of reasons as such.
- Meaning matters in philosophy because the specifically conceptual contents of declarative sentences consist in the role they play in reasoning.

Adopting a pragmatics-first order of explanation, we pursued this reason-centered approach by distinguishing discursive practices as those in which some performances have the significance of assertions and denials, expressions of commitment to accept or reject claimables. That

significance we understand as depending on practices of defending and challenging claimings, giving reasons for and against them, thereby affecting practitioners' entitlements to their commitments. This pragmatic discursive rationalist account gave us our first glimpse of and grip on doxastic reason relations of implication and incompatibility among claimables. Along the way from there we have seen how to implement and fill in the details of our further strategic commitments, to understand reason relations and conceptual contents in accordance with the principles of logical expressivism, semantic inferentialism, and bimodal conceptual realism.

The result is an intricately structured quartet of rational metaconceptual vocabularies, that is, vocabularies for specifying reason relations and conceptual contents: a bilateral normative pragmatic vocabulary, two kinds of model-theoretic semantic vocabulary (one representationalist and one inferentialist), and a particularly expressively powerful but otherwise familiar logical vocabulary presented in a sequent-calculus format. It is worth rehearsing some of the relations among the members of this quartet of vocabularies for explicitly expressing reason relations and conceptual contents. To begin with, we offered reasons to think that they represent something like a complete set of kinds of rational metavocabularies. The bilateral pragmatic metavocabulary and the truth-maker semantic metavocabulary are extrinsic to and explanatory of the reason relations of the base vocabularies they address, while the logical and implication-space vocabularies are intrinsic to and explicative of the reason relations of the base vocabularies from which they are elaborated. Within those two broad kinds, the pragmatic and semantic metavocabularies address the practices of using declarative sentences of ground-level vocabularies and their representational meanings, respectively. Logical vocabulary expresses the reason relations of arbitrary base vocabularies by extending those very vocabularies, while the implication-space metavocabulary codifies those reason relations and the conceptual contents they articulate in a new metavocabulary built out of raw materials supplied by the base, but mentioning rather than using the expressions of that base vocabulary.

We have shown how metaconceptual vocabularies of all these four general kinds can be so constructed that they have the expressive power to make explicit the reason relations (and the rational forms or conceptual contents they articulate) not only of structurally closed base vocabularies—the kind the philosophical tradition by and large restricts itself to considering—but also of those that have structurally open or substructural nonmonotonic (and even hypernonmonotonic) and nontransitive reason relations. At the center of this achievement is the logic NMMS. There are many formulations of classical logic in the proof-theoretic vocabulary of the sequent calculus that are all equivalent in the

context of the full topological structure that Gentzen and Tarski both discern in and impose on logical consequence relations. In substructural settings, these specifications come apart and behave differently. We exploit that diversity to find the expressive sweet spot occupied by NMMS. Ketonen's invertible connective rules, with judicious mixing of additive and multiplicative rules (usually avoided in order to secure the admissibility of the structural transitivity principle of Cut), remain well-behaved and exhibit hitherto unprecedented expressive power when we use them to extend not only structurally closed but also more relaxed open-structured base vocabularies. The purely logical reason relations of the logically extended vocabularies that result from base vocabularies with nonmonotonic and nontransitive reason relations are wholly classical—and, indeed, structurally closed: monotonic and transitive. In this sense we offer not a nonmonotonic (or, more generally, substructural) logic, but a logic for expressing nonmonotonic (and, more generally, substructural) material implications and incompatibilities.

Following the clues provided by this bilateral pragmatics for base vocabularies with open-structured reason relations, we showed how to define consequence in the truth-maker semantic metavocabulary so that it is sound and complete for our logic—and thereby how to extend the truth-maker framework to codify open-structured, nonmonotonic and nontransitive material reason relations. The implication-space metavocabulary also supplies a sound and complete semantics for NMMS, and for open-structured reason relations generally. Following that clue, we showed how to construct a thorough-going isomorphism between the extrinsic-explanatory representational truth-maker semantic metavocabulary and the intrinsic-explicative implication-space semantic metavocabulary. Any constellation of reason relations, open or closed, that can be expressed by the one kind of semantic metavocabulary can be expressed by the other. This is a deep connection between representationalist model-theoretic semantic frameworks and inferentialist model-theoretic semantic frameworks. In addition, the interpretation we offered of multiplicative and additive linear logic in implication-space semantics made visible a fundamental homology between the basic operations on conceptual roles in that semantics (adjunction and symjunction, premisory and conclusory roles) and basic features of sequent-calculus proofs (such as single-premise-sequent and multi-premise-sequent connective rules, left and right rules)—a structural connection between the proof theoretic sequent-calculus metavocabulary in which we construct our logic and the implication-space model-theoretic metavocabulary that we take to be the intrinsic semantic metavocabulary of reason relations.

The right kind of logical vocabulary, NMMS, accordingly ties together the other three kinds of rational metavocabulary. The bilateral normative pragmatic metavocabulary on which our more deontically fine-grained version is based was introduced in the first place to explain the implication relations expressed as sequents in multisuccedent proof-theoretic sequent calculi for various logics. Both the truth-maker semantics and the implication-space semantics are sound and complete for that logic, and seeing why and how they are reveals an isomorphism between them at the level of reason relations. The final tie that binds together the quartet of rational metavocabularies is the isomorphism between the bilateral pragmatic metavocabulary and the truth-maker semantic metavocabulary when just the right conception of consequence is defined in the latter framework. A premise-set Γ implies a conclusion-set Δ in the deontic pragmatic sense if and only if commitment to accept all of Γ precludes entitlement to deny all of Δ, and in the alethic semantic sense if and only if every fusion of truth-makers of all of Γ with false-makers of all of Δ is an impossible state. Working out this detailed isomorphism between reason relations specified in terms of norms governing subjects' use of sentences and what, in virtue of that very isomorphism, we can see as the same reason relations specified in terms of the worldly representational meaning of those sentences (what they say about how things objectively are) is offering an account of the relations between language and the world. In our account, the representational language-world nexus is to be understood at the deepest level in terms of reason relations, and only later, and as a consequence, at the level of sentences and what makes them true or false. That account makes visible the sense in which rational forms, articulated by relations of consequence and incompatibility, can be shared by what people say and the world they are talking about. The implication-space semantic metavocabulary makes those rational forms explicit. It is their native language.

In closing let us offer one final perspective on this constellation of rational metavocabularies. The discursive rationalism that motivates and shapes our enterprise understands vocabularies in general as the medium of consciousness. "Consciousness" in this sense is conceptual awareness, Kantian apperception: sapience rather than the mere sentience that consists in the capacity to have raw feels such as pains and sensations of red. Being conscious in this sense is being able to accept or reject, assert or deny, say, think, believe or suppose that p, where p is some declarative sentence. Consciousness so understood is a matter of responding to things by applying concepts, making claims. The claimable, conceptual contents expressed by using declarative sentences to assert and deny, and to challenge and defend assertions and denials, are what they are because of the reason relations of implication and incompatibility they

stand in to other such conceptual contents. Every set of such reason relations is accordingly a form of rational consciousness. The four kinds of vocabulary we have considered for making explicit in different ways the reason relations of arbitrary base vocabularies are accordingly forms of rational *self*-consciousness: rational consciousness of what rational consciousness is. Each of these rational metavocabularies adopts and expresses a distinctive perspective on reason relations: pragmatic, logical, representational-semantic, and inferential-semantic. The strategy of metaconceptual functionalism is to understand reason relations as whatever one can understand in these four intricately interrelated ways. By considering the relations among all of them, the final functional account of reason relations is a higher, more complete form of rational self-consciousness than any of the four kinds of rational metavocabulary provides on its own. Metaconceptual rationalism as filled in here accordingly shows up as an account of the essential internal structure of rational self-consciousness as such.

Notes

1 Our motivation for this is, of course, related to a variant of Frege's context principle. Hence, we don't think of this as a culpable omission but rather as laying the foundations for an investigation of subsentential expressions. A defense of this claim is, however, beyond the scope of this book.

2 As a reminder of just one of the many rare virtues that our theories have, recall, for instance, that our disjunction obeys so-called "disjunction simplification," which says that if $\Gamma, A \vee B \mathrel{|\!\sim} \Delta$, then $\Gamma, A \mathrel{|\!\sim} \Delta$ and $\Gamma, B \mathrel{|\!\sim} \Delta$. And as Fine (1975) pointed out, no theory of nonmonotonic consequence that allows intersubstitutions of sentences that are classically equivalent can do that.

3 Here we followed Daniel Kaplan, who showed in his PhD dissertation (Kaplan, 2022), how one can generalize Girard's phase-semantics for linear logic so as to allow for material implications and failures of Cut.

Epilogue
A Speculative Synthesis

Ulf Hlobil

The implication-space semantics that we presented in Chapter Five marked the endpoint of our journey from the phenomenon of norms governing combinations of assertions and denials and the phenomenon of compossibility of worldly states to a realm of abstract entities that are the roles that can be identified in both of these phenomena. What we want to do in this epilogue is not an integral part of our accounts of reason relations, content, rational forms, or logic. It is rather a highly speculative and optional reflection on these accounts.[1] These speculations are inspired by the traditions of German idealism and (some versions of) scholastic Aristotelianism, although our ideas differ in important ways from those that can be found in either of these traditions.

Readers with no taste for this kind of speculative philosophy can ignore this epilogue without missing anything that is essential to the foregoing chapters. However, what we say below strikes us as an illuminating—though optional—perspective on our project. Since our project has many new and perhaps radical aspects, appreciating this perspective might be helpful, even for a reader who is unsympathetic to this style of philosophy.

In medieval terms, one arrives at a proposition *a posteriori* if one arrives at it on the basis of an appreciation of its consequences or effects. Kant calls this the "analytic method" (*Jäsche Logic*, §117, AA9, 149). In contrast to this, one arrives at a proposition *a priori* (in the medieval sense) if one arrives at it on the basis of an appreciation of prior, more fundamental, or higher principles or causes. Kant calls this the "synthetic method." The way in which we have arrived at implication-space semantics was by distilling its structure out of two approaches to reason relations and, hence, conceptual content. The pragmatics-first approach uses a pragmatic-normative metavocabulary to give an account of reason relations. And the semantics-first approach uses a semantic-representationalist metavocabulary to give such an account. If these two approaches are—in one sense of that term—less fundamental than implication-space semantics, then we have arrived at implication-space semantics *a posteriori*

or by following the analytic method. In this epilogue, we want to sketch an attempt to pursue our project *a priori*, or, in Kantian terms, we want to pursue a synthetic methodology.

We will start with the most abstract characterization of rational forms—that is, propositional conceptual contents—from Chapter Five, and our aim is to explain how this yields the two incarnations of this structure with which we started. In particular, we will start with the monoidal structure of reason relations that we explained in Section 5.3. In the first section of this epilogue, we will render this monoidal structure in philosophical terms, thereby putting some philosophical content on the merely formal structure. In the second section, we combine this philosophical rendering of the fundamental structure of reason relations with what we call the "Rational Conflict Assumption" to arrive at the notion of a world that is the standard of accuracy and truth for occurrences of rational forms. In the third section, we want to show how the notions of discursive acts governed by a broadly normative kind of modality and worldly states governed by a broadly alethic kind of modality emerge from the notion of a world in a synthetic way, given some assumptions that we will highlight along the way. We will thus arrive at the idea of worldly states governed by a broadly alethic modality that we have encountered in Chapter Four. And we will end by returning to the pragmatic ideas of discursive acts of making claims that are governed by a broadly normative modality with which we began this volume in Chapters One to Three.

Rational Forms as Starting Points

As we saw in Chapter Five, we can think of rational forms as roles that bearers of these rational forms can play in reason relations. In particular, we saw that sentences have a particular rational form in virtue of their assertions and denials playing a particular role with respect to further assertions and denials. And worldly propositions have a particular rational form in virtue of their truth-makers and falsity-makers playing a particular role with respect to further states. What it means to play such roles is for the assertions and denials or the truth-makers and falsity-makers to be incompatible with constellations of further assertions and denials or further states, respectively. The incompatibility among assertions and denials is normative, whereas the incompatibility among worldly states is alethic. These incompatibility relations encode the reason relations that are our central topic.

Recall from Section 5.3 that the essential mathematical structure of reason relations is the structure of a commutative monoid that is defined on a set of pairs, together with a bipartition of that set of pairs. Thus, the basic elements of this structure are pairs of bearers of rational forms:

candidate implications. There is a commutative and associative operation (with an identity element) of combining such implications to yield further implications. And a bipartition divides all implications into the excluded ones—the one that are out-of-bounds or impossible—and the ones that are not excluded.

This structure carries over to rational forms, which are the roles that bearers of rational forms can play in the monoid of reason relations just sketched. First, rational forms are implicational roles that are pairs of a premisory role and a conclusory role. Thus, we can say that rational forms have two sides or poles, namely their premisory role and their conclusory role. Second, rational forms can be combined to yield further rational forms. In particular they can be combined by the operation of adjunction, which is associative and commutative (and has an identity element). Third, there is a partition between rational forms for which all implications in the adjunction of their premisory and conclusory roles are excluded (so that every implication with a premise that has the premisory role of the rational form and a conclusion that has the conclusory role of the rational form is a good implication, see the definition of content entailment in Chapter Five) and those rational forms for which this is not the case. So, just like reason relations, so also rational forms have the structure of a commutative monoid that is defined on a set of pairs, together with a bipartition of the monoid set. This structure is the most abstract and general characterization of rational forms. For it to be more than merely a mathematical structure, we must spell out the notion of exclusion that defines the bipartition and the notion of combination that underlies the monoidal operation. By spelling out the philosophical significance of these notions, we can start to unpack the abstract mathematical structure of rational forms.

We will now sketch a synthetic and philosophical development of this abstract and merely formal characterization of rational forms. We will assume that we do not know what worldly states or discursive acts like assertions and denials are, and we will assume that we do not know which bearers of rational forms are governed by what kind of modality. Our aim is to let an understanding of these facts emerge from a development and unpacking of the fundamental nature of rational forms. Accordingly, we begin by spelling out in philosophical terms the fundamental features of rational forms that we have just sketched in formal terms.

The first feature of rational forms is that they have two sides or poles, which we may call their positive and negative poles. In the formal structure, this feature corresponds to the fact that the elements of our monoid set are pairs, where the two sides or poles of rational forms are the two elements of these pairs. To understand the philosophical significance of this, it is helpful to notice that we can consider rational forms at three levels. At the first level, a repeatable, abstract entity can have a rational form. (We

have seen in the analytic telling of our story that these entities will turn out to be sentences and worldly propositions.) At the second level, these abstract entities can occur in two different ways, which we call "positive" and "negative" ways to occur. Every occurrence of a rational form is an occurrence either of its positive pole or of its negative pole. In the first case the rational role occurs in a positive way, and in the second case the role occurs in a negative way. The positive and negative ways for a rational form to occur correspond to their premisory and conclusory roles in the implication-space semantics from Chapter Five. The occurrence of a rational form is a determinable, and its only two determinates are the positive and the negative occurrence of the rational form. At a third level, a particular object or event or fact is a particular positive or negative occurrence of one (or more) rational form(s). (We have seen in the analytic telling of our story that particular positive occurrences of rational forms—occurrences of their premisory roles—can be interpreted as assertions or as the obtaining of truth-makers, while particular negative occurrences of rational forms—occurrences of their conclusory roles—can be interpreted as denials or the obtaining of falsity-makers.) It is only at the third level that rational forms occur enmattered in particulars, and they always do so by way of the particular being a positive or negative occurrence of the rational form. We can summarize the situation as follows:

Level 1	rational form
Level 2	poles of the rational form: positive or negative
Level 3	particular positive or negative occurrences of the rational form

The philosophical significance of the bipolarity of rational forms is that every particular occurrence of a rational form is an occurrence of either its negative or its positive pole. Rational forms and their poles are repeatables, but their individual occurrences are particulars.

A second essential feature of rational forms is that particular occurrences of positive and negative poles of various rational forms can be combined with each other, which corresponds to the operation of adjunction in the implication-space semantics. In this way, many occurrences of the positive and negative poles of different rational forms can form a unity: their combination can be enmattered, actualized, or particularized together. We can model such unities in implication-space semantics by implications; for implications are combinations of positive occurrences of the rational forms of their premises and negative occurrences of the rational forms of their conclusions. Thus, we can think of the combination of positive and negative poles of rational forms as itself being the occurrence of a rational form,

and we can think of the rational forms that occur in it as the forms whose positive and negative poles are combined in the larger occurrence.

For particular occurrences of rational forms, we can then think of the largest particular occurrence of a rational form of which the particular occurrence is a part as the unified constellation of occurrences of rational forms in which the rational forms are enmattered together. So, for each particular occurrence of a rational form, there is a particular unified constellation in which it occurs. There are no particular occurrences of rational forms that are not included in exactly one unified constellation of occurrences of rational forms. We call such a unified constellation to which occurrences of rational forms belong a "medium" for rational forms. Thus, a medium is a unified constellation in which poles of rational forms can occur or be enmattered together. (We have seen in the analytic telling of our story, that these media will turn out to be the subjects that hold positions in a discursive practice, and a world in which states obtain together. So, we will work towards an understanding of media as the positions of subjects in a discursive practice, on one side, and the world, on the other side.) If an occurrence of a rational form is enmattered in a particular medium, we say that the medium includes the occurrence of the rational form. A constellation of occurrences of rational forms that are not actually enmattered together but could be enmattered together is a merely possible medium. When the constellation of occurrences of rational forms that are a possible medium are enmattered together, this is an actual medium.

A third essential feature of rational forms is that combinations of their occurrences can be excluded or not. This corresponds to the partition of the space of implications into the excluded ones (the good ones) and all others (the bad ones). To put it differently, some occurrences of rational forms are compatible with each other and others are not. As a matter of terminology, we say that two occurrences, or combinations of occurrences, are incompatible or mutually exclusive if and only if their combination is excluded. Negative and positive occurrences of the same rational form are always incompatible. This is the sense in which the two poles of a rational form are opposed to each other, and this corresponds to the structural feature of Containment in the formal theories from previous chapters. (We have already seen in the analytic telling of our story that the nature of the exclusion relation can be normative or alethic, and we will return to this below. One of our goals is to understand this bifurcation between two kinds of modality from a synthetic perspective.) Notice that the relevant mutual exclusion relations can hold between particular (combinations of) positive or negative occurrences of rational forms within a single medium; this is the intra-medium aspect of incompatibility. However, there is also an inter-medium aspect of the incompatibility relation. We can formulate the difference thus:

Intra-medium-conflict: Incompatible occurrences of rational forms are present in one medium.

Inter-medium-conflict: Incompatible occurrences of rational forms are present in different media.

These two kinds of conflict between occurrences of rational forms are of the same generic kind: they are kinds of incompatible (positive or negative) occurrences of rational forms. They differ merely in whether the incompatible occurrences of rational forms are present in a single medium or are distributed over several media. For simplicity, we will usually think of merely two incompatible occurrences of rational forms and of at most two media in inter-medium-conflicts. In a case of inter-medium conflict, we will say that the two media disagree because they include incompatible occurrences of rational forms.[2]

We want to add a fourth and final essential feature of rational forms, namely that they are individuated by the incompatibility relations in which their positive and negative occurrences stand. That is, if the positive and negative occurrences of some rational form stand in all and only the incompatibility relations that the positive and negative occurrences of "another" rational form stand in, then these rational forms are identical. In other words, if there are two distinct rational forms, then there is some positive or negative occurrence of one of these rational forms that is compatible with something with which the positive or negative occurrence of the other rational form is not compatible. This is modal functionalism about rational forms: rational forms are individuated by the roles that their occurrences play in incompatibility relations. In implication-space semantics, this fourth essential feature of rational forms comes out in the fact that rational forms are, in effect, defined as equivalence classes with respect to exclusion relations (mediated by the notion of ranges of subjunctive robustness).

If rational forms are as we just described them, then we can use the following characterization as the starting point for our synthetic methodology:

Divided Unities There are rational forms and their essence is that they are divided unities, in the sense that:

(a) every rational form can occur in two different ways (poles), positively or negatively;

(b) such occurrences can be combined into larger unities, which occur together in media;

(c) the occurrences can (and some do) stand in incompatibility relations to each other (for instance, occurrences of the two poles of any rational form are incompatible); and

(d) rational forms are individuated by the incompatibility relations of their positive and negative occurrences.

It is worth noting that since incompatibility relations encode reason relations, clause (d) of Divided Unities implies that rational forms are individuated by the reason relations in which their occurrences stand. In other words, rational forms are roles that bearers of rational forms can play in reason relations: they are, as we said in Chapter Five, implicational roles or conceptual contents.

In the remainder of this epilogue, we want to sketch a way in which one may try to unfold and develop what is implicit in Divided Unities. Taking Divided Unities as our first principle, we will pursue a synthetic path back to the starting points of our investigation, namely the truth-maker theory of Chapter Four and the normative pragmatics of Chapters One to Three.

Ideal Occurrences of Rational Forms: The World

We begin with clause (c) of Divided Unities, namely that positive and negative occurrences of rational forms can stand in incompatibility relations. In this section, we want to show that if we add an assumption, which we will call the "Rational Conflict Assumption," to Divided Unities, then we can introduce a distinction between media that are like subjects who accept and reject claims and an ideal medium that can be understood as the world that these subjects represent, in the sense of answering to it for the correctness of their claims. We can thus arrive at the idea of the subject-object nexus. One way to understand our project in this section is as an attempt to answer the question: What do we need to add to Divided Unities in order to introduce the distinction between discursive acts of subjects, on one side, and worldly states of objects, on the other side? Our answer will be that we need add only the Rational Conflict Assumption. In the context of Divided Unities, the Rational Conflict Assumption is sufficient for the subject-object nexus to emerge, and it may or may not also be necessary.

We have said above that a medium is that within which a constellation of particular occurrences of rational forms is combined into a unity. And this is the force of clause (b) of Divided Unities.[3] When there are two media in which rational forms occur, then the two poles of one rational form may occur, one in each of these two media. In this case, the rational form occurs

in mutually exclusive ways in the two media, as is ensured by clause (c) of Divided Unities, which says that the two poles of the same rational form are always mutually exclusive. More generally, the occurrences of rational forms in two media may be incompatible, even when they are not the two poles of a single rational form. These are cases of inter-medium conflicts.

So far, our notion of incompatibility—that is, our notion of an exclusion of combinations of occurrences of rational forms—is merely a partition of all possible combinations of occurrences of rational forms into two sets, the excluded ones and all the other combinations. We have not said anything about what this partition means. What does it mean for occurrences of rational forms to be excluded? Or, equivalently, what does it mean for occurrences of rational forms to be incompatible? By using terms like "incompatible," Divided Unities suggests that there is some kind of conflict or tension between incompatible occurrences of forms. How can we understand this tension or conflict?

An inter-medium conflict is, at a minimum, a disagreement, in the sense that neither of the two disagreeing media can include the incompatible occurrence of rational forms from the other medium, on pain of including within itself incompatible occurrences of rational forms. If the first medium includes a positive occurrence of a rational form, for instance, and the second medium includes a negative occurrence of the same rational form, then adding a negative occurrence of the rational form to the first medium yields a situation in which the combination of all occurrences of rational forms in the first medium is excluded. Let us call this "noncotenability." Noncotenability is a rather weak notion of disagreement. In order to see this, let us distinguish it from two other kinds of disagreement[4]:

1. *Noncotenability*: One medium includes an occurrence of a positive or negative pole of a rational form and a second medium includes an occurrence of a positive or negative pole of a (possibly different) rational form, and the combination of the occurrence in the first medium and the occurrence in the second medium are incompatible.
2. *Proselytizing conflict*: From the perspective of each medium, the incompatible occurrence of rational forms in the other medium is incorrect.
3. *Rational conflict*: At most one—and sometimes exactly one—of the two incompatible occurrences of rational forms in the two disagreeing media is correct, in the sense that any positive (negative) occurrence of this rational form in any medium is correct if and only if the occurrence that is correct in the disagreement at issue is positive (negative).

As we have seen, the kind of inter-medium disagreement that we are considering gives rise to noncotenability. One may wonder, however,

whether there must also be inter-medium disagreements that are proselytizing conflicts or even rational conflicts. The notions of correctness that are operative in the definitions of proselytizing conflicts and rational conflicts are notions of the property of the occurrence of a rational form in virtue of which the conflict holds and that is the occurrence in the medium in favor of which the conflict is to be decided, whatever that property may turn out to be.

Here we want to introduce a strong and substantive assumption. The rest of this epilogue should be understood as conditional on this assumption. The assumption is that there are rational forms that underwrite the existence of rational conflicts between some possible media.

Rational Conflict Assumption
There are rational conflicts between possible media.

Let us use the term "objective rational form" for a rational form that is such that some possible media stand in a rational conflict because one medium includes a positive occurrence and the other medium includes a negative occurrence of the rational form at issue *and* exactly one of these occurrences of the form is correct.

The Rational Conflict Assumption is a strong assumption. If it is true, then there are objective rational forms, that is, rational form such that either the positive or the negative occurrence is correct, and this assessment of correctness carries over to any occurrence of that rational form in any other medium. It is the same standard of correctness that applies to all media that disagree about any given objective rational form. In this sense, the Rational Conflict Assumption implies that there is an objective standard of correctness for all occurrences of objective rational forms.[5]

The Rational Conflict Assumption allows us to introduce the notion of an ideal medium. This is the idea of a medium in which all and only objective rational forms occur, and all occurrences of objective rational forms in the medium are correct, that is, they agree with the objective standard of correctness. Thus, for each objective rational form the ideal medium includes either a positive or a negative occurrence and not both. The kind of occurrence (positive or negative) that is included in the ideal medium is the correct one. The ideal medium is the embodiment, in a medium, of the standard of correctness for occurrences of rational forms. It is the medium in which all objective rational conflicts are decided correctly; that is, the correct pole and only the correct pole of each objective rational form occurs in the ideal medium.

The ideal medium is unique, in the sense that although there could be several particular media that meet the definition of the ideal medium, they would all include all and only the same occurrences of the same poles of all

objective rational forms. Therefore, the ideal medium is unique as a kind of medium, or as a possible medium. The combination of occurrences of rational forms in the ideal medium is itself an occurrence of a rational form, by the combination mentioned in clause (b) of Divided Unities, and it is the same kind of occurrence of the same rational forms in all media that are ideal media, by clause (d) of Divided Unities. Thus, there is a unique kind of occurrence of a rational form that is the ideal medium (although there might be several particular media that instantiate this kind of occurrence). So, we are justified to talk about *the* ideal medium (as a kind of medium or a possible medium).

It is irrelevant to our present discussion whether anyone could, even in principle, decide a disagreement between media by determining what occurrences of rational forms are included in the ideal medium.

Since the ideal medium includes all and only the occurrences of objective rational forms that are correct, whether an occurrence of an objective rational form in any nonideal medium is correct varies with whether it matches the polarity of the occurrence of this rational form in the ideal medium. We can think of this matching as accuracy[6] and say: A positive occurrence of a rational form is *accurate* just in case the rational form occurs positively in the ideal medium. And a negative occurrence of a rational form is *accurate* just in case the rational form occurs negatively in the ideal medium. Thus, accuracy is a correspondence between the polarity of the occurrence of an objective rational form in some medium and the polarity of its occurrence in the ideal medium.

We can now introduce the notions of truth and falsity of rational forms in terms of the accuracy of positive and negative occurrences of these rational forms. While accuracy is a property of occurrences of poles of rational forms, truth is a property of rational forms themselves. We say that a rational form is *true* if and only if it is an objective rational form whose positive occurrence is accurate. And we say that a rational form is *false* if and only if it is an objective rational form whose negative occurrence is accurate.[7] To put it differently, true rational forms are those that occur positively in the ideal medium, and false rational forms are those that occur negatively in the ideal medium. Since rational forms are identical to propositional conceptual contents, we have accordingly introduced a notion of truth and falsity for propositional conceptual contents. Moreover, positive occurrences of true rational forms in nonideal media are true because the form occurs positively in the ideal medium, and occurrences of false rational forms in nonideal media are false because the form occurs negatively in the ideal medium. In this sense, the ideal medium is the standard of truth and falsity. A rational form that is true is true if and because it occurs positively in the ideal medium. And a rational form that is false is false if and because it occurs negatively in the ideal

medium. If what is represented is that which determines the truth or falsity of representations, then the ideal medium is what is represented by occurrences of objective rational forms in nonideal media. (This part of the synthetic telling of our story corresponds, in the analytic telling in the body of this work, to the discussion of representation at the start of Chapter Four.)

Let us use "the world" for the standard of truth and falsity, such that what is true is true because of what the world is like and what is false is false because of what the world is like. If what we said so far is correct, the world is the ideal medium. According to this way of using the term "world," the world is the ideal medium in which all objective rational forms occur. So, the world is complete or maximal with respect to objective rational forms: every rational form occurs in the world either positively or negatively. And the world is consistent in the sense that the combination of occurrences of objective rational forms that occur in the world is not excluded; it is coherent. And if any objective rational form were to occur in the world in a way (positively or negatively) in which it does not already occur in the world, this would render the world incoherent (because it would require that the world includes a positive and also a negative occurrence of the same rational form). So, the world is a maximal coherent combination of occurrences of objective rational forms.[8]

The world is articulated by reason relations. That is, it is articulated by incompatibility relations that individuate the positive and negative poles of rational forms that occur in it. And we can also express this by saying that conceptual contents are enmattered in the world. (This is the conceptual realism that we endorsed in the Introduction and in Chapter Four. If we combine this with what we already know from the analytic telling of the story, we can say: What is and is not the case in the world is the same as what can be asserted or denied, what can be true or false. In the synthetic order of explanation, however, discursive acts and worldly states are introduced only in the next section.)

Let us say that a rational form "is the case" when it occurs positively in the world. So what can be the case or not are rational forms, and rational forms are the case if and only if they are true. It might seem that what is true and what makes something true are identical, thus yielding an identity theory of truth. However, truth is not an intrinsic feature of a rational form. Rather, what makes the rational form true is that it occurs positively in the world. So, what is the case and what is true are identical, but what makes something true is not what is the case (the rational form) but that it is the case (the positive occurrence of the rational form in the world). In this sense, truth is correspondence with what is the case.

To sum up, given Divided Unities and the Rational Conflict Assumption, we can arrive at the ideas of the world and of truth as correspondence with

the world via the idea of rational conflicts and objective rational forms. We started with the notion of rational forms provided by Divided Unities, and the Rational Conflict Assumption then implies that some rational forms are such that either their positive or their negative occurrences (and not both) are correct for any medium. The world is the kind of medium in which the correct poles of all and only these rational forms occur. Occurrences of rational forms in nonideal media are accurate just in case the world includes an occurrence of the same rational form with the same polarity (positive or negative). A rational form is true if and only if it occurs positively in the world.

The Pragmatic-Semantic Bifurcation Explained

In the previous section, we have seen how Divided Unities and the Rational Conflict Assumption jointly yield the notions of a world and of truth as correspondence with what is the case in the world. Our topic in this section is how a bifurcation arises between the two approaches to content that we have used as our guides in previous chapters. Our goal is to chart a synthetic path to the distinction between discursive acts that are governed by a broadly normative kind of modality and worldly states that are governed by a broadly alethic kind of modality. According to the picture that will emerge below, the pragmatics-first approach to content aims to explain, in a deontic-normative metavocabulary, the exclusions among the occurrences of rational forms as discursive acts in nonideal media. The semantics-first approach to content aims to explain, in an alethic-modal metavocabulary, the exclusions among occurrences of rational forms as worldly states in the ideal medium, which is the world.

To see how the bifurcation between the two approaches arises in a synthetic way from the forgoing developments, it is helpful to notice two points. First, the assessment of occurrences of rational forms in the world as accurate is trivial. There is, by the very nature of the world as the ideal medium, no room for inaccuracy in the occurrences of rational forms in the world. Thus, although the world can serve as the standard for normative assessments in terms of accuracy of other media, such assessments do not apply to the world in any nontrivial way. So, an understanding of incompatibility relations as normative relations whose ultimate standard of assessment is agreement with the world does not apply (nontrivially) to the world itself.

Second, the world does not include any occurrences of rational forms that are incompatible with each other (within the context of the other occurrences of rational forms in the world). For, if the occurrences of rational forms in virtue of which two media disagree could all agree with the ideal medium (against the background of the other occurrences of

rational forms in the ideal medium), then agreement with the ideal medium could not serve as a standard for deciding the disagreement between the disagreeing media.[9] To see this, suppose for reductio that the ideal medium can include occurrences of objective rational forms that are mutually incompatible. Then two media could disagree in virtue of the first including a particular (positive or negative) occurrence of some objective rational form and the second one including a particular (positive or negative) occurrence of some objective rational form, and they could nevertheless both agree with the ideal medium. Hence, the ideal medium could not serve as the standard by which the rational disagreement between the two media is decided, which contradicts the definition of the ideal medium. So all the occurrences of rational forms in the ideal medium are compatible with each other; the combination of occurrences of rational forms in the ideal medium is not excluded. If we try to think of incompatibility relations as normative in a sense such that a medium including incompatible occurrences of rational forms constitutes some kind of impropriety, the world cannot be guilty of such improprieties. The world cannot violate this normative standard.

We can distinguish two broad genera of incompatibility relations, namely incompatibility relations such that it is possible for incompatible relata of the relation to be actual and those for which this is not the case. We call the first kind of incompatibility relations "broadly normative" incompatibility relations and the second kind "broadly alethic" incompatibility relations.[10] Broadly normative incompatibility relations allow for improprieties in the sense of allowing for the actuality of incompatible items. Broadly alethic incompatibility relations do not allow for such improprieties. In this sense, normative incompatibility relations can be violated and alethic incompatibility relations cannot be violated. We can summarize this distinction as follows.

Genera of Incompatibility
An incompatibility relation is broadly alethic if and only if it is impossible that incompatible relata of this relation are actual, and it is broadly normative otherwise.

The two points we have made above imply that the incompatibility relations that hold among potential occurrences of rational forms in the ideal medium—the world—are of a broadly alethic and not a broadly normative kind. For, the world does not include mutually incompatible occurrences of rational forms.

It makes sense to call the kind of incompatibility that we find in the world broadly alethic. For, if the combination of the positive occurrences of two forms is excluded, then the two rational forms cannot be true

together. If the combination of the negative occurrences of two forms is excluded, then the two rational forms cannot be false together. And if the negative occurrence of one form excludes the positive occurrence of another, it cannot be the case that the first is false and the second true. Thus the incompatibility relation between positive (or negative) occurrences of rational forms manifests itself in the world as the noncompossibility of the truth (or falsity) of these rational forms. In virtue of the world being the ideal medium, the excluded combinations of occurrences of rational forms cannot occur in the world.

Let us turn to the incompatibility relations among occurrences of rational forms in nonideal media. The exclusion relations between occurrences of rational forms in nonideal media are of the broadly normative and not the broadly alethic kind. For, nonideal media can include actual occurrences of rational forms that are mutually incompatible. Nothing in what we said prevents, for instance, a nonideal medium that includes a positive and also a negative occurrence of a single rational form. And such occurrences of rational forms are incompatible—indeed, they are incompatible against any background of further occurrences of rational forms.[11] Similarly, different nonideal media can include occurrences of rational forms that are incompatible. When this happens, at least one of the occurrences of a rational form is incorrect or mistaken, according to the standard provided by the ideal medium. It is characteristic of normative assessments to underwrite the possibility of errors. When two things are alethically incompatible, however, they cannot both be actual. Therefore, the incompatibility relations between occurrences of rational forms in a single nonideal medium and also incompatibilities between different media are not broadly alethic but broadly normative incompatibilities.

The upshot of what we just said is that the incompatibility relations in the world are broadly alethic-modal relations and the incompatibility relations within and among nonideal media are broadly normative-modal relations. These incompatibility relations among rational forms (determining assessments of the coherence of constellations of occurrences of those forms) are, however, isomorphic. Indeed, by clause (d) of Divided Unities, they must be isomorphic in order for the same rational forms to occur in the different media. (This isomorphism was the main result of Chapter Four, in the analytic telling of our story.)

Let us call occurrences of rational forms in nonideal media "discursive acts." And let us use "acceptance" and "rejection" for the positive and negative ways, respectively, in which rational forms occur in nonideal media. Furthermore, let us call occurrences of rational forms in the ideal medium of the world "states." And let us use "truth-maker" and "falsity-maker" for the positive and negative ways, respectively, in which rational forms occur in the ideal medium, that is, in the world. So every discursive

act is an acceptance or rejection of a rational form (and perhaps also an acceptance and rejection of other rational forms as well), and every state is a truth-maker or a falsity-maker of a rational form (and perhaps also a truth-maker or a falsity-maker of other rational forms as well).

We have now arrived at the bifurcation between discursive acts that are governed by broadly normative incompatibility relations, and worldly states that are governed by broadly alethic incompatibility relations. Media in which rational forms occur in discursive acts are positions or views of subjects (nonideal media), and the medium in which rational forms occur in worldly states is the world (the ideal medium). Thus, there are several nonideal media in which rational forms occur in discursive acts governed by normative incompatibilities, and one ideal kind of medium in which rational forms occur in worldly states governed by alethic incompatibility relations.

The normative bilateralist interpretation of sequent calculi offers a vocabulary for specifying which combinations of acceptances and rejections of rational forms are normatively excluded. (This was the topic of Chapter Three, in the analytic telling of our story.) The modal bilateralist interpretation of truth-maker semantics offers a vocabulary for specifying which combinations of truth-makers and falsity-makers are alethic-modally excluded. (This was the topic of Chapter Four, in the analytic telling of our story.) We can now see the first as specifying a theory of occurrences of rational forms in nonideal media, while we can see the second as specifying a theory of occurrences of rational forms in the ideal medium. The isomorphism between the two theories is the isomorphism that must hold for representation and inter-medium disagreement to be possible.

Finally, we can say that in nonideal media, what it means for the combination of a positive occurrence of a rational form and a negative occurrence of another rational form to be excluded is that if a nonideal medium (a discursive subject) is committed to accepting the first rational form, then the nonideal medium cannot be entitled to reject the second rational form. And what it means for a combination of positive occurrences of two rational forms to be excluded is that a nonideal medium (a discursive subject) cannot be entitled to accept both rational forms. In the first case, the first rational form is a reason for the second. And in the second case, the first rational form is a reason against the second. These are the reason relations of being a reason for and being a reason against that, according to the normative pragmatics of Chapter One, constrain practices of giving and asking for reasons. We have accordingly reached the end point of the synthetic telling of our story, which was the starting point of the analytic telling of our story.

This concludes our sketch of a synthetic path connecting the main ideas of this book. Let us sum up. If, in the context of the structure of rational forms we have called "Divided Unities" we make the Rational Conflict Assumption, then we can define the notion of an ideal medium, which embodies the standard of correctness that makes the most demanding kind of rational conflicts possible. The ideal medium is that in virtue of which what we assert or deny is true or false; it is the world. Since the world, in virtue of being the ideal medium, cannot include mutually incompatible occurrences of rational forms, the exclusion relations among potential occurrences of rational forms in the world cannot be broadly normative exclusion relations, but must be broadly alethic exclusion relations. Normative standards require the possibility of errors or failures to accord with those standards. In contrast to this, the exclusion relation among occurrences of rational forms in nonideal media are broadly normative exclusion relations because nonideal media can include incompatible occurrences of rational forms. We can accordingly divide media for the occurrences of rational forms into two broad functional classes, depending on the general features of the notion of incompatibility they exhibit. Those two broad functional classes are the modal genera comprising alethic modal incompatibility and normative modality (or at least, genera whose paradigmatic species are those two kinds). These are the more specific structures characteristic of the two sides of the intentional nexus: the subjective and the objective. These are the species we began by investigating, the first in normative pragmatic metavocabularies for reason relations, the second in truth-maker semantic metavocabularies for reason relations.

Conclusion

This highly speculative epilogue sketched a line of thought that is not an essential part of the project of this book. Our aim in this epilogue was not to convince skeptical readers of the viability of thinking about the world as the ideal medium for occurrences of rational forms, or of the possibility of deducing the subject-object nexus from the fundamental structure of rational forms. Rather, our aim was to illustrate what a synthetic telling of the story that we present in the rest of this book might look like.

If one wants to tell this story in a way that follows a synthetic methodology, one must start from the highest, most abstract principles. We suggested that Divided Unities might be such a principle. To start with Divided Unities is, in effect, a way to begin our story with the abstract structure of implication-space semantics. Following a synthetic methodology, one could then try to unfold and develop what is implicitly contained in the most abstract principles. We suggested that Divided

Unities, when combined with the Rational Conflict Assumption, implicitly contains the ideas of an ideal medium that is the world, and the division between alethic and normative kinds of incompatibility relations as the ways in which the incompatibility relations mentioned in Divided Unities manifest themselves in the ideal medium and in nonideal media, respectively. To put it differently, by unfolding and developing the notions of incompatibility and unity, a bifurcation emerges between two kinds of media in which rational forms can be enmattered: an ideal medium in which rational forms are enmattered in worldly states and nonideal media in which rational forms are enmattered in discursive acts. This bifurcation then gives rise to a distinction between two kinds of incompatibility relations, between two kinds of modality: alethic and normative modality. In this way, we came back to the initial starting point of this book: the pragmatics of normative bilateralism.

In this synthetic telling of our story, we had to assume that the notions of incompatibility and unity that occur in Divided Unities (and, less directly, in the implication-space semantics from Chapter Five) are not barren formal notions but rather implicitly contain fruitful philosophical ideas. Moreover, we relied on the Rational Conflict Assumption. We have used this assumption in order to illustrate what a synthetic telling of our story might look like, and we do not claim that it is true. We provide this sketch of a potential synthetic telling of our story because it seems to us that it is difficult fully to appreciate many of the details of the analytic telling of our story without having at least a rough idea of what a synthetic telling of that story might look like. We hope that this epilogue provides such a sketch, despite its speculative, optional, and tentative nature.

Notes

1 We disagree regarding the plausibility of the Rational Conflict Assumption that we use in this epilogue. While Hlobil thinks that this assumption can be justified (although we do not attempt to do so here), Brandom is more skeptical.

2 We suppress some details about the formal structure of rational forms. For instance, the formal fact that the monoidal operation is associative and commutative reflects the assumption that the order and grouping of the combination of occurrences of rational forms does not matter. And that there is an identity element allows us to encode whether a combination of occurrences of rational forms is excluded in the range of subjunctive robustness of that combination. Thus, when we define rational forms in terms of ranges of subjunctive robustness, it becomes essential to a rational form whether its occurrences are excluded. We will mostly suppress such details in the remainder of this chapter.

3 From our analytic telling of the story, we already know that the kind of unity that media have will be the unity of positions of discursive subjects, on one side, and the unity of the world, on the other side. Kantians might want to explain

both kinds of unity in terms of the synthetic unity of apperception. We here merely presuppose that media are unified wholes without offering any account of what unifies them.

4 This classification is inspired by, but different from, that offered by MacFarlane (2014), who distinguishes practical and doxastic noncotenability, preclusion of joint accuracy, and preclusion of joint reflexive accuracy. Since we are only interested in occurrences of rational forms and do not specify these further, the difference between practical and doxastic noncotenability does not apply. Our proselytizing conflict corresponds roughly to MacFarlane's preclusion of joint accuracy. And our rational conflict corresponds roughly to MacFarlane's preclusion of joint reflexive accuracy. MacFarlane's notions are defined using technical terms that we do not want to take on board.

5 One thing to mean by "anti-realism," in the current context, is the thesis that there are no rational conflicts between media. Anti-realism in this sense is, at an abstract level, the kind of view that Huw Price (2003) characterizes as admitting a subjective assertability norm (sincerity) and a norm of personal warranted assertability (justification) but not a third norm of truth. Translating these norms into our setting, we could say that the sincerity norm says that media should not display occurrences of rational forms that do not actually occur in them, and the justification norm says that media should not include the occurrence of a positive (negative) pole of a rational form if they do not also include occurrences of rational forms that are incompatible with the opposite negative (positive) pole of the first rational form. The truth norm would say that a medium should include an occurrence of an objective rational form only if it meets the objective standard of correctness, that is, only if it would agree with the correct side in a rational conflict.

Price writes: "The important point is that this [third norm] provides a norm of assertion which we take it that a speaker may fail to meet, even if she does meet norms of subjective assertibility and (personal) warranted assertibility. We are prepared to make the judgment that a speaker is incorrect, or mistaken, in this sense, simply on the basis that we are prepared to make a contrary assertion; in advance, in other words, of any judgment that she fails to meet one or other of the two weaker norms" (Price, 2003, 176). Price argues that we need the third norm in order to make possible a discursive practice with the right kind of friction. Price's arguments for the necessity of the third norm and, hence, against anti-realism strike us as supporting the Rational Conflict Assumption—although perhaps not decisively.

6 In order to avoid misunderstandings, we should look ahead and point out that we use "accuracy" to talk about the correctness of the poles of rational forms, and we use "truth" to talk about the correctness of rational forms themselves. This differs from how the term "accuracy" is sometimes used by philosophers. Sometimes accuracy is treated as gradable, in contrast to truth. Sometimes it is meant to apply to states of creatures that do not possess concepts. Neither of these contrasts is part of our notion of accuracy.

7 The way in which we introduce the notions of truth and falsity here is closely related to the way in which Restall (2009) introduces truth-values. Restall first

introduces the notion of a limit position, which is a lot like an invalid ("out of bounds") sequent in which every sentence occurs either on the left or on the right. Restall, in effect, says that we can think of truth (relative to a limit position) as the status that a sentence has when it occurs on the left in a limit position. And falsity (relative to a limit position) is the status that a sentence has when it occurs on the right in a limit position. To see the similarity to our notions of truth and falsity, notice that one can think of the ideal medium as characterized by an invalid sequent in which, for every objective rational form, a sentence that expresses that rational form occurs either on the left or on the right side, that is, the sentence occurs either as a premise (positively) or as a conclusion (negatively). Thus, the ideal medium plays a role that is very similar to Restall's limit position. Moreover, we say that a rational form is true if it occurs positively in the ideal medium, and it is false if it occurs negatively in the ideal medium. In contrast to Restall, however, we have not imposed structural principles like Weakening or Cut. And there may be rational forms (namely those that are not objective rational forms) that occur neither positively nor negatively in the ideal medium. Nevertheless, it may be helpful to view our strategy here as similar to what Restall does.

8 In this respect, the world as an ideal medium is similar to worlds in possible world semantics, where worlds are sometimes understood as maximal consistent sets of propositions.

9 There can be disagreements between media such that there is noncotenability but there is no rational conflict because the incompatibility between the occurrences of rational forms in virtue of which the two media disagree is cured by the addition of other occurrences of rational forms in the ideal medium. In this sense, there may be no rational conflict even though the occurrences of rational forms that are incompatible in two media are the kind of occurrences of rational form that can stand in rational conflicts. This happens when the disagreement is merely apparent because the occurrences of rational forms do not exclude each other against the background of the occurrences of rational forms in the ideal medium. Thus there are, merely apparent disagreements; and there might not be any way for us to find out whether a disagreement is merely apparent, in this sense, or genuine. The converse is also possible. That is, there may be occurrences of rational forms that are compatible relative to the other occurrences in their respective media but are incompatible relative to the occurrences of rational forms in the ideal medium. We may call these "masked rational disagreements." And again there might be no way for us to find out whether an apparent agreement is really a masked rational disagreement. These possibilities are consequences of the fact that we allow reason relations to be open reason relations and, hence, nonmonotonic. Such complications are ignored in the body of this chapter.

10 In terms familiar from modal logic, our distinction corresponds to modalities for which the modal axiom called "M" or "T" (which says: $\Box\phi \rightarrow \phi$) holds and those for which it fails.

11 This follows from clause (c) of Divided Unities, and it corresponds to the formal principle of Containment in the setting of the current discussion.

Bibliography

Angell, R. B. (1977). Three systems of first degree entailment. *Journal of Symbolic Logic*, 42:147.

Angell, R. B. (1989). Deducibility, entailment and analytic containment. In Norman, J. and Sylvan, R., editors, *Directions in Relevant Logic*, volume 42, pages 119–143. Kluwer Academic Publishers, Dordrecht.

Aristotle (2014a). *Complete Works of Aristotle: Volume 1 (revised Oxford translation, edited by J. Barnes)*. Princeton University Press, Princeton.

Aristotle (2014b). *Complete Works of Aristotle: Volume 2 (revised Oxford translation, edited by J. Barnes)*. Princeton University Press, Princeton.

Armstrong, D. M. (2004). *Truth and Truthmakers*. Cambridge University Press, Cambridge.

Barrio, E. A., Rosenblatt, L., and Tajer, D. (2015). The logics of strict-tolerant logic. *Journal of Philosophical Logic*, 44(5):551–571.

Beall, J. and Murzi, J. (2013). Two flavors of Curry's paradox. *Journal of Philosophy*, 110(3):143–165.

Belnap, N. (1962). Tonk, Plonk and Plink. *Analysis*, 22(6):130–134.

Berto, F. and Restall, G. (2019). Negation on the Australian plan. *Journal of Philosophical Logic*, 48(6):1119–1144.

Bimbó, K. (2015). *Proof Theory: Sequent Calculi and Related Formalisms*. CRC Press, Boca Raton.

Brandom, R. B. (1994). *Making It Explicit: Reasoning, Representing, and Discursive Commitment*. Harvard University Press, Cambridge, Mass.

Brandom, R. B. (2000). *Articulating Reasons: An Introduction to Inferentialism*. Harvard University Press, Cambridge, Mass.

Brandom, R. B. (2008). *Between Saying and Doing: Towards an Analytic Pragmatism*. Oxford University Press, Oxford.

Brandom, R. B. (2019). *A Spirit of Trust: A Reading of Hegel's* Phenomenology. Harvard University Press, Harvard.

Brower, J. E. and Brower-Toland, S. (2008). Aquinas on mental representation: Concepts and intentionality. *Philosophical Review*, 117(2):193–243.

Bueno, O. and Shalkowski, S. A. (2013). Logical constants: A modalist approach. *Noûs*, 47(1):1–24.

Cobreros, P., Egré, P., Ripley, D., and van Rooij, R. (2012). Tolerant, classical, strict. *Journal of Philosophical Logic*, 41(2):347–385.

Cobreros, P., Égré, P., Ripley, D., and van Rooij, R. (2013). Reaching transparent truth. *Mind*, 122(488):841–866.

Cobreros, P., Egré, P., Ripley, D., and van Rooij, R. (2020a). Inferences and metainferences in ST. *Journal of Philosophical Logic*, 49(6):1057–1077.

Cobreros, P., Rosa, E. L., and Tranchini, L. (2020b). (I can't get no) Antisatisfaction. *Synthese*, 198(9):8251–8265.

Correia, F. (2016). On the logic of factual equivalence. *Review of Symbolic Logic*, 9(1):103–122.

Correia, F. and Skiles, A. (2019). Grounding, essence, and identity. *Philosophy and Phenomenological Research*, 98(3):642–670.

Cross, C. B. (2003). Nonmonotonic inconsistency. *Artificial Intelligence*, 149(2):161–178.

deRosset, L. and Fine, K. (2022). A semantics for the impure logic of ground. *Journal of Philosophical Logic*, 52:415–493.

Dicher, B. and Paoli, F. (2019). ST, LP and tolerant metainferences. In Bakent, C. and Ferguson, T. M., editors, *Graham Priest on Dialetheism and Paraconsistency*, pages 383–407. Springer, Cham.

Dretske, F. (1981). *Knowledge and the Flow of Information*. The MIT Press, Cambridge, Mass.

Dummett, M. (1973). *Frege: Philosophy of Language*. Harper and Row, New York.

Elgin, S. Z. (2021). The semantic foundations of philosophical analysis. *Review of Symbolic Logic*, 16(2):603–623.

Fine, K. (1975). Critical notice of Lewis, Counterfactuals. *Mind*, 84(335):451–458.

Fine, K. (2014). Truth-maker semantics for intuitionistic logic. *Journal of Philosophical Logic*, 43(2–3):549–577.

Fine, K. (2016). Angellic content. *Journal of Philosophical Logic*, 45(2):199–226.

Fine, K. (2017a). A theory of truthmaker content I: Conjunction, disjunction and negation. *Journal of Philosophical Logic*, 46(6):625–674.

Fine, K. (2017b). A theory of truthmaker content II: Subject-matter, common content, remainder and ground. *Journal of Philosophical Logic*, 46(6):675–702.

Fine, K. (2017c). Truthmaker semantics. In Hale, B., Wright, C., and Miller, A., editors, *A Companion to the Philosophy of Language*, volume 2, pages 556–577. Wiley Blackwell, New York.

Fine, K. (2018a). Compliance and command I: Categorical imperatives. *Review of Symbolic Logic*, 11(4):609–633.

Fine, K. (2018b). Compliance and command II: Imperatives and deontics. *Review of Symbolic Logic*, 11(4):634–664.

Fine, K. (2018c). Truthmaking and the is-ought gap. *Synthese*, 198(2):887–914.

Fine, K. (2019). Verisimilitude and truthmaking. *Erkenntnis*, 86(5):1239–1276.

Fine, K. (2020). Yablo on subject-matter. *Philosophical Studies*, 177(1):129–171.

Fine, K. and Jago, M. (2019). Logic for exact entailment. *Review of Symbolic Logic*, 12(3):536–556.

Fitting, M. (2021). A family of strict/tolerant logics. *Journal of Philosophical Logic*, 50(2):363–394.

Frege, G. (1956). The thought: A logical inquiry. *Mind*, 65(259):289–311.

Frege, G. (1979). *Posthumous Writings*. Basil Blackwell, Oxford.

Frege, G. (1998). *Begriffsschrift und andere Aufsätze*. G. Olms, Hildesheim.

French, R. (2016). Structural reflexivity and the paradoxes of self-reference. *Ergo: An Open Access Journal of Philosophy*, 3:113–131.

Gabbay, D. M. (1985). Theoretical foundations for nonmonotonic reasoning in expert systems. In Apt, K., editor, *Logics and Models of Concurrent Systems*, pages 439–459. Springer, Berlin and New York.

Gentzen, G. (1934). Untersuchungen über das logische Schließen: I. *Mathematische Zeitschrift*, 39(2):176–210.

Girard, J.-Y. (1987). Linear logic. *Theoretical Computer Science*, 50(1):1–101.

Girard, J.-Y. (2001). Locus solum: From the rules of logic to the logic of rules. *Mathematical Structures in Computer Science*, 11(3):301–506.

Hale, B. (2020). Exact truthmakers, modality, and essence. In Leech, J., editor, *Essence and Existence: Selected Essays*, pages 124–140. Oxford University Press, Oxford.

Harman, G. (1984). Logic and reasoning. *Synthese*, 60(1):107–127.

Harman, G. H. (1986). *Change in View: Principles of Reasoning*. MIT Press, Cambridge, Mass.

Hempel, C. G. (1963). Implications of Carnap's work for the philosophy of science. In Schilpp, P., editor, *The Philosophy of Rudolf Carnap*, pages 685–709. Open Court, La Salle.

Hlobil, U. (2016). A nonmonotonic sequent calculus for inferentialist expressivists. In Arazim, P. and Danák, M., editors, *The Logica Yearbook 2015*, pages 87–105. College Publications, London.

Hlobil, U. (2017). When structural principles hold merely locally. In Arazim, P. and Lávika, T., editors, *The Logica Yearbook 2016*, pages 53–67. College Publications, London.

Hlobil, U. (2018). Choosing your nonmonotonic logic: A shoppers' guide. In Arazim, P. and Lávika, T., editors, *The Logica Yearbook 2017*, pages 109–123. College Publications, London.

Hlobil, U. (2020). Expressing validity: Towards a self-sufficient inferentialism. In Blicha, M. and Sedlár, I., editors, *The Logica Yearbook 2019*, pages 67–82. College Publications, London.

Hlobil, U. (2022a). The laws of thought and the laws of truth as two sides of one coin. *Journal of Philosophical Logic*, 52:313–343.

Hlobil, U. (2022b). A truth-maker semantics for ST: refusing to climb the strict/tolerant hierarchy. *Synthese*, 200(5):1–23.

Horty, J. (2007). Defaults with priorities. *Journal of Philosophical Logic*, 36(4):367–413.

Humberstone, L. (2011). *The Connectives*. MIT Press, Cambridge, Mass.

Jago, M. (2018). *What Truth Is*. Oxford University Press, Oxford.

Kant, I. (1992). *Lectures on Logic*. The Cambridge edition of the works of Immanuel Kant. Cambridge University Press, New York.

Kaplan, D. (2018). A multi-succedent sequent calculus for logical expressivists. In Arazim, P. and Lávička, T., editors, *The Logica Yearbook 2017*, pages 139–153. College Publications, London.

Kaplan, D. (2022). *Substructural Content*. PhD thesis, University of Pittsburgh, USA; online at: http://d-scholarship.pitt.edu/42065/.

Kraus, S., Lehmann, D., and Magidor, M. (1990). Nonmonotonic reasoning, preferential models and cumulative logics. *Artificial Intelligence*, 44:167–207.

Lakatos, I. (1976). *Proofs and Refutations: The Logic of Mathematical Discovery*. Cambridge University Press, Cambridge.

MacFarlane, J. (2014). *Assessment Sensitivity: Relative Truth and its Applications*. Clarendon Press, Oxford.

McDowell, J. H. (1996). *Mind and World*. Harvard University Press, Cambridge, Mass.

Moltmann, F. (2020). Truthmaker semantics for natural language: Attitude verbs, modals, and intensional transitive verbs. *Theoretical Linguistics*, 3:159–200.

Nair, S. (2019). Must good reasoning satisfy cumulative transitivity? *Philosophy and Phenomenological Research*, 98(1):123–146.

Negri, S., Von Plato, J., and Ranta, A. (2008). *Structural Proof Theory*. Cambridge University Press, Cambridge.

Pasnau, R. (2004). Form, substance, and mechanism. *Philosophical Review*, 113(1):31–88.

Price, H. (1990). Why 'not'? *Mind*, 99(394):221–238.

Price, H. (2003). Truth as convenient friction. *Journal of Philosophy*, 100(4):167–190.

Priest, G. (2006). *In Contradiction: A Study of the Transconsistent*. Oxford University Press, Oxford.

Ré, B. D., Szmuc, D., and Teijeiro, P. (2021). Derivability and metainferential validity. *Journal of Philosophical Logic*, 51(6):1521–1547.

Reiter, R. (1980). A logic for default reasoning. *Artificial Intelligence*, 13(1–2):81–132.

Restall, G. (2005). Multiple conclusions. In Hájek, P., Valdés-Villanueva, L., and Westerstaahl, D., editors, *Logic, Methodology and Philosophy of Science*, pages 189–205. College Publications, London.

Restall, G. (2009). Truth values and proof theory. *Studia Logica*, 92(2):241–264.

Restall, G. (2013). Assertion, denial and non-classical theories. In Berto, F., Mares, E., Tanaka, K., and Paoli, F., editors, *Paraconsistency: Logic and Applications*, pages 81–99. Springer, Dordrecht.

Ripley, D. (2012). Conservatively extending classical logic with transparent truth. *Review of Symbolic Logic*, 5(2):354–378.

Ripley, D. (2013). Paradoxes and failures of Cut. *Australasian Journal of Philosophy*, 91(1):139–164.

Ripley, D. (2015). Anything goes. *Topoi*, 34(1):25–36.

Rumfitt, I. (2000). Yes and no. *Mind*, 109(436):781–823.

Sagi, G. (2018). Logicality and meaning. *Review of Symbolic Logic*, 11(1):133–159.

Santorio, P. (2018). Alternatives and truthmakers in conditional semantics. *Journal of Philosophy*, 115(10):513–549.

Sellars, W. (1963). Abstract entities. *Review of Metaphysics*, 16:627–671.

Sher, G. (1991). *The Bounds of Logic: A Generalized Viewpoint*. MIT Press, Cambridge, Mass.

Shimamura, S. (2017). A nonmonotonic modal relevant sequent calculus. In Baltag, A., Seligman, J., and Yamada, T., editors, *Logic, Rationality, and Interaction. LORI 2017. Lecture Notes in Computer Science*, pages 570–584. Springer, Berlin.

Shimamura, S. (2019). A first-order sequent calculus for logical inferentialists and expressivists. In Sedlár, I. and Blicha, M., editors, *The Logica Yearbook 2018*, pages 211–228. College Publications, London.

Simonelli, R. (2022). Considering the exceptions: On the failure of cumulative transitivity for indicative conditionals. *Synthese*, 200(5):1–21.

Simonelli, R. (2023). Why must incompatibility be symmetric? *Philosophical Quarterly*, pqad078.

Smiley, T. (1996). Rejection. *Analysis*, 56(1):1–9.

Spinoza, B. (2002). *Spinoza: Complete Works*. Hackett Publishing Company, Indianapolis.

Steinberger, F. (2011). Why conclusions should remain single. *Journal of Philosophical Logic*, 40(3):333–355.

Takeuti, G. (1987). *Proof Theory*. North-Holland, Amsterdam, 2 edition.

Tanter, K. (2021). Subatomic inferences: An inferentialist semantics for atomics, predicates, and names. *Review of Symbolic Logic*, 16(3):672–699.

Tarski, A. (1936). Über den Begriff der logischen Folgerung. In *Actes du Congrès International de Philosophie Scientifique, VII*, pages 1–11. Hermann et Cie, Paris.

van Heijenoort, J. (1967). *From Frege to Gödel: A Source Book in Mathematical Logic, 1879–1931*. Harvard University Press, Cambridge, Mass.

Wittgenstein, L. (1953). *Philosophical Investigations*. Blackwell, Oxford.

Yablo, S. (2014). *Aboutness*. Princeton University Press, Princeton.

Yablo, S. (2016). Ifs, ands, and buts: An incremental truthmaker semantics for indicative conditionals. *Analytic Philosophy*, 57(1):175–213.

Zylstra, J. (2019). Making semantics for essence. *Inquiry: An Interdisciplinary Journal of Philosophy*, 62(8):859–876.

Index

acceptance 30–8, 42–51, 55, 104, 283–7, 292, 321–2; Pragmatically Implicit (PIA) 47; *see also* assertion; practical attitudes; speech acts
alethic 6–9, 11–12, 165–9, 181, 197–99, 206n13, 290–4, 309, 319–24; impossibility 30–2, 165–6, 170–2, 182–4, 285; necessity 285, 291; noncompossibility 320–21; *see also* impossibility; modality; necessity
Antecedent-Adjunction (AA) 115, 125, 129, 146, 301
assertion 3, 10–11, 106, 156n10, 159–70, 180–4, 206n12–13, 210, 306–9, 325n5; as expressing acceptance 6, 37, 42–5, 51–3, 55–7, 284; *see also* acceptance; commitment; denial; speech acts
axioms: of classical logic 72, 151–2; of NMMS 104–5, 108–10, 112, 118, 134, 136–8, 141, 204, 246–7, 300; *see also* base; sequent calculus

base 5–7, 12–20, 22–8, 61–71, 75–8, 80–2, 84, 86–7, 92–7, 105, 107–8, 110–121, 126–31, 136–49, 151–7, 190–1, 193–200, 211, 221–2, 224, 227, 231,235–6, 254, 259, 262, 266–70, 282, 285–6, 288, 294, 297, 300, 302, 304–7; autonomous 26, 105, 107, 111–2, 126, 131, 196; substructural 16–17, 72, 80–1, 89, 96–7, 100, 148–9, 195, 197, 255, 297, 304–5; *see also* expressivism; LX; vocabulary
Bayesianism 99
bilateralism 10–12, 24–5, 42–54, 98–9, 196–8, 293, 304–6; normative 32, 47–8, 56–7, 106–7, 182, 192, 221–4; truth-maker 179, 182, 192, 224–8; *see also* bivalence; isomorphism; metavocabulary; Restall
bivalence 53, 55–6, 143; semantic 29–32, 41; *see also* bilateralism; truth-values
Brandom, Robert B. 19, 40, 44, 101n2, 101n4, 206n14, 324n1

Carnap, Rudolf 1, 10, 59
Cautious Monotonicity (CM) 82–91, 93, 139–42, 148, 296; *see also* closure; explicitation; hypernonmonotonicity
challenge 6, 37, 43–5, 50–3, 55–6, 72–4, 144–5; *see also* claiming; defense, incompatibility; justification; speech acts
claimable 3, 8–9, 35–8, 42–6, 56, 61, 77, 83–4, 94, 96, 145, 304; as object of practical attitudes 30–2, 35, 69; *see also* conceptual content; propositions
claiming 3, 9, 26, 37, 40, 42–4, 182, 196; asserting 6, 68, 172; denying 38, 89, 156n10; as expressing practical attitudes 30–2, 35–8, 44–6, 55–6, 69; *see also* assertion; denial; discursive practices; practical attitudes
classical logic xi, 17, 23, 30, 65, 72, 75, 80–2, 96, 99, 101n8, 104, 106, 112–15, 118–21, 136–42, 148–9, 151, 153–6, 173–4, 189–90, 193–4,

224, 242, 244, 300–3, 305, 307n2; axioms of 106, 118–21; classicality operator 136–7; and Containment 17, 82, 96, 112–5, 118–21, 193; and NMMS 17, 80–1, 112–15, 151, 153–6, 224, 300–3; structure of 72, 80–2, 101n8, 104, 112–5, 118–21, 242, 305; *see also* logicism; NonMonotonic MultiSuccedent logic (NMMS)
closure 16, 23, 60, 70, 72, 76 , 78, 80–2, 85, 87–92, 95–9, 101n8, 103,147–9, 160, 173, 182, 296, 303–5; explicitation 85, 87–92, 95–9, 103, 296, 303–5; operator 72; structural 23, 60, 70, 72, 76, 101n8, 103, 182, 187, 304–5; topological 60, 72, 76, 78, 80–2, 103, 296; *see also* classical logic; consequence; monotonicity; open; reflexivity; transitivity
CM *see* Cautious Monotonicity
CO *see* Containment
commitment 8–12, 24, 38, 40, 43–54, 56–7, 89–90, 92, 101n1, 106, 124–5, 144, 197, 282, 296, 304, 306; doxastic 8–12, 43–54, 144, 197, 282, 296, 304, 306; implicit 47–8, 56–7, 89–90, 92, 101n1, 124–5, 296; *see also* acceptance; deontic; entitlement; rejection
completeness 27, 156n11, 189, 200, 202, 222, 231, 254, 263; expressive 69, 146
conceptual connections 68, 249–53, 279n23
conceptual content 1–2, 5, 8–9, 18, 22–4, 26–8, 49, 53, 56–9, 61–4, 70, 86, 94, 99–100, 105, 145, 157n19, 164, 199–200, 211, 215–21, 224, 226–33, 249–50, 276n7–8, 282, 284–97, 301, 303–7; as functionally specified 26, 53, 145, 211, 282; and implicational roles 215–17, 226–8, 281, 285; as inferential role 4, 8–9, 18, 22–4, 49, 56, 61, 99–100, 145, 199–200, 211, 281, 294–5, 303–7; propositional/of sentences 4–5, 22–4, 57, 105; *see also* meaning; practical attitudes; propositions; rational forms; reason relations
conceptual realism 8–11, 159, 167, 318; and rational forms 9, 163–7, 198, 290–2
conceptual role 4, 17–28, 56, 99, 105, 199–200, 210–76, 291, 294; *see also* implicational role; implication-space semantics; rational form; role inclusion
conditional xii, 15–16, 27, 34–6, 63–4, 67–8, 73–7, 79–80, 100, 123, 130–1, 139–40, 146, 156n10, 219, 222–3, 297, 301; expressive role of 34–6, 63–4, 67–8, 95, 100, 115, 123, 139–40, 146, 301; and implication 63–4, 100, 146; subjunctive 34–5, 73–5, 291; *see also* Deduction-Detachment; implication; logical expressivism; sentential connectives
conic models 242–247, 269–272, 299, 303
conjunction 16, 65, 68, 95, 103, 113–5, 146, 175–6, 185–6, 200, 207n20, 219, 232–3, 241, 246, 252, 261, 265–7, 301
connective rules 14–15, 66, 70, 95, 144–7, 305; *see also* expressivism; sequent calculus; structural rules
consciousness, rational self- 25, 28, 69, 122–123, 126, 294, 306–7
conservativeness 15–17, 64–6, 80–1, 95–6, 111, 114–5, 118, 139–40, 148; *see also* elaboration; tonk
consequence xi, 3–6, 8–11, 16–17, 24–8, 30, 39, 41–2, 47–8, 52–3, 57–60, 62–3, 72, 74, 77–98, 101–2, 104–8, 111–6, 118–21, 127–30, 132–6, 139–42, 147–9, 154–5, 178–82, 185–95, 197–8, 202, 204, 208–9, 210, 220–2, 224–7, 231, 233, 235–6, 241–6, 254–5, 259–62, 264, 269, 272–3, 279, 291–3, 296, 300–3, 305–7; classical (CL) 41–2, 72, 104–8; logical 16–17, 30, 47–8, 52–3, 57–60, 220–2, 241–6, 259–62, 264, 269, 272–3; material 16, 24–28, 39, 47–8, 62–3; Pragmatic-Normative (PN) 10–11, 46–9, 210; substructural 72, 74, 77–98, 101–2, 111–6, 118–21, 132–6, 139–42, 147–9, 154–5;

Truth-Maker 178–82, 185–95, 197–8, 224–7; *see also* base; bilateralism; implication; incompatibility; reason relations; truth-maker consequence
Containment (CO) 15, 17, 68, 72, 81–2, 85, 87, 95–6, 108, 110–13, 115, 118–19, 135–7, 143, 147, 151–2, 154, 178, 180, 182, 187–90, 193–4, 201–2, 207, 209, 224, 239, 243–4, 278–9, 298, 300–1, 303, 312, 326; *see also* closure; sequent calculus; structural rules
Contraction 96, 105, 110–1, 129–31, 138–9, 142, 148–9, 150, 155n3, 157n25, 207n20, 209n34, 230–2, 254, 262, 279n19, 296, 302–3; *see also* idempotence
contradiction 68, 78, 93, 97, 155n2, 242; *see also* incompatibility; negation
contradictory (Aristotelian) 29, 68; *see also* incompatibility; negation
contrariety (Aristotelian) 45, 68; *see also* incompatibility; negation; reasons
Correia's logic 246–8, 274–5, 280n26, 289, 299–300, 303; *see also* metaphysics
Covariant Tracking 163, 168, 209n32, 277n8
CT *see* Cumulative Transitivity
Cumulative Transitivity (CT) 79–80, 83–5, 87–91, 97, 101n10–1, 119–120, 139, 141–2, 296; *see also* closure; Cut; hypernonmonotonicity; structural rules
Cut 41, 83, 86, 101n3, 101n11, 137, 146–147, 155n2, 157n26, 180, 182, 187–92, 201–2, 207n15, 235–7, 239, 242, 244–5, 278n16, 279, 298, 303, 305, 307n3; Mixed-Cut 110–1, 116–7, 119, 151–2; PN 182; soundness of 189; TM 182, 188–9, 192, 201–2; *see also* sequent calculus; structural rules

Davidson, Donald 2, 207n16
DD *see* Deduction-Detachment
Deduction-Detachment (DD) 67–8, 76–7, 79, 109, 115, 123, 125, 129, 139–40, 146, 219, 222–3, 301; *see also* conditional; explicitation; expressivism
defeasible/defeasor 72–4, 77–8, 80, 86, 93
defense 6, 45, 52–3, 55, 74, 144; *see also* challenge; claiming; implication; justification; reasoning; reasons; reasons relations; speech acts
denial 8–11, 30, 32–3, 35, 37–8, 42–5, 50, 52, 55–6, 89–90, 106, 122, 124–5, 159–70, 172, 180–4, 192, 195–6, 206n12–3, 210, 288, 304, 307–11, 323; *see also* negation; practical attitudes; rejection; speech acts
deontic 7, 9–12, 30–2, 291–2; commitment/entitlement 9–10, 12, 40, 44–6, 52–3, 56, 306; *see also* commitment; entitlement; modality
derivability 39, 53, 67, 105, 109–12, 115, 146, 187, 302
Descartes, René 11, 59
Dewey, John 2
discursive practices 2–3, 6–10, 26, 31, 35, 38–40, 45, 52–3, 55–7, 92, 100, 104, 107, 143, 196–9, 285, 290–4, 304, 309–10, 314, 319, 321–2, 325n5; claiming 8, 26, 31, 69, 92, 100, 167, 196, 304, 309; critical character 3, 35, 45, 52–3, 55, 92, 100, 104, 143, 304, 325n5; *see also* asserting; challenge; claiming; defense; justification; pragmatics; speech acts
disjunction 16, 65, 68, 103, 113–5, 146, 176–7, 186–7, 200, 207n20, 232–3, 236, 241, 246, 258, 266–7, 301, 307n2
distinctions: acceptance/rejection 30–5, 45–8, 55–6, 206n13, 285–9, 292–3, 321–2; adjunction/symjunction 216–8, 223, 257, 297–8; alethic/deontic modality 7–11, 30–2, 291–2, 319–24; assertion/denial 10–11, 52, 55–7, 159–61, 165–8, 181; Australian/American plan for negation 68; challenge/defense 45, 52–3, 55–6; CM/CT 83–91, 139; commitment/entitlement 40, 42–52,

56–7, 106, 197–8, 304; connective/structural rules 128–31, 147, 184–9, 230–7; contractive/noncontractive 146–7, 230–7, 262–3, 278n12, 302–3; expressive/aggregative sentential connectives 15, 68; expressivism/logicism 57–71, 78–82, 86–9, 94–5, 104, 121–4, 143; explicit/implicit 48–9, 56–7, 64–5, 83–9, 91–4, 102n14, 121–4; extrinsic-explanatory/intrinsic-explicative metavocabularies 12–13, 17, 24, 27–8, 199–200, 281–2, 286, 304–5; form/matter 163–7, 169, 205n9, 210, 290–4, 311–2; implication/incompatibility 4–9, 12–16, 25–6, 37–43, 46–50, 53–4, 55–6, 66–7, 77–8, 97–8, 196–8, 276n2, 282–3, 285; inferentialism/representationalism 18–9, 22–5, 255, 294, 304–5; model theory/proof theory 18, 23–5, 107–15, 180–4, 221–4, 235–7, 298; nonmonotonic/hypernonmonotonic 85, 95–6; open/closed structure 72–3, 89, 91, 187, 296, 304–5; reasons for/against 3–6, 26, 35–40, 42–3, 55–6, 74, 145, 293; *salva veritate/salva consequentia* 101n6, 215–6; semantics/pragmatics 2–4, 29–32, 195–7
Downward-Closure 177, 192–4, 208n23; and soundness of Weakening 188–9, 201; *see also* truth-maker semantics; Weakening
Dummett, Michael 2, 57–8

elaboration 14–17, 64, 66–7, 70–1, 94–6, 107, 144–5, 300, 302
entitlement 9, 40, 44–53, 56, 106, 196–7, 304; default 45, 53, 56; preclusion of 9, 11, 47–8, 50, 56–7, 106, 282, 306
equivalence class: and implicational roles 215, 226, 295; of sentences 33, 34, 218
equivalence relation 27, 98; among implicational roles 212–4, 276n1, 276n6, 277n8, 286
essence 173, 247–8, 255–6, 272, 274, 300; *see also* grounding
ex contradictione quodlibet 4, 25, 78, 97, 102n15, 242, 244; *see also* implicational roles; negation
exclusion 284–5, 288–9, 323; logical 211, 221, 228; mathematical 310; modal 195, 294; normative 291, 312, 321; rational 4, 25, 30, 56; *see also* incompatibility
Exclusivity 177, 180, 183, 188–91, 194, 201–2, 204, 209n34, 209n40, 298; *see also* Containment; Strong Kleene Logic (K3)
Exhaustivity 177, 180, 182, 188–94, 201–2; and soundness of Cut 189, 201; *see also* Cut; Strong Kleene Logic (K3)
explicitation/explication xi, 69, 84–93, 98, 101n12, 102n15, 121–9, 139, 145–6, 296; by implication 84–93, 121–9, 131, 145–6, 152; of implicit content 84–93, 98, 102n15; of reason relations 121–9; by sequents 69, 126–9; *see also* inference; logical expressivism; logical metavocabulary; logical vocabulary; logicism, LX
explosion *see ex contradictione quodlibet*
expressivism xi, 17, 55, 60–4, 69–70, 80, 86, 94, 101, 103–5, 107, 111, 114, 121–3, 126, 128, 131, 133, 139–41, 143–5, 147–8, 199, 235, 282, 294, 301–2, 304; *see also* elaboration; explicitation; logicism; LX; NonMonotonic MultiSuccedent logic (NMMS)
extension 14, 64, 81, 97, 118, 121, 141–2, 151, 170, 190, 288; conservative 23, 114, 139, 148, 297

fact 9, 34–5, 160–1, 163, 168, 171, 206–7, 234–5, 247–8, 252, 255, 264–5, 272, 274, 289–90, 297, 299, 311
falsifier *see* falsity-maker
falsity 11, 29–32, 35, 53, 55, 159–63, 165–6, 168–9, 172, 174–81, 184–9, 191, 195, 203, 205–7, 210, 225–6, 245, 248, 252, 260, 272–5, 284–5, 287, 289, 293–4, 309, 311, 317–18, 321–2, 325–6; *see also* bivalence; truth; truth-values

falsity-maker 160–3, 165–6, 168–9, 174–81, 184–9, 191, 195, 203, 205–7, 210, 225–6, 245, 248, 252, 260, 272–5, 284–5, 287, 293–4, 309, 311, 321–2; *see also* representation; truth conditions; truth-maker; truth-maker semantics
Fine, Kit 1–2, 10, 19–21, 32, 161, 173–4, 177–8, 182, 189, 205n8, 207n17, 207n19, 208n22, 208n23, 208n25 246, 307n2; *see also* truth-maker semantics
form *see* rational form; hylomophism
Frege, Gottlob 1, 9, 20, 27–9, 34, 57–62, 70–1, 143–4, 149, 307
functionalism 9, 11–12, 26, 28, 281, 285; metaconceptual 290, 307; metalinguistic 9, 26, 28, 285, 290; meta-rational 200; in philosophy of mind 93; about rational forms 164–167, 170, 313; about reasons 98; about reason relations 99, 290
fusion 7, 11, 24, 174–6, 178–9, 180–1, 185–9, 194, 197–8, 225, 306

Gentzen, Gerhard 13, 18, 25, 41–2, 50, 72, 83, 87–8, 95–7, 102n14, 105–106, 145–6, 149, 151–2; *see also* Cut; sequent calculus
Girard, Jean-Yves xi, 96, 234–5, 276n1, 278n14, 283, 297, 307n2
grounding 173, 255–6, 272, 300; strict full 248, 275; *see also* essence; metaphysics

Harman, Gilbert 36–38, 43, 46, 48–9, 89–90, 104
Heidegger, Martin 2
Hempel, Carl G. 60
hylomorphism 9, 163–5, 167, 291; compounds 163, 165, 167; matter 163–7, 169, 205n9, 210, 253, 285–7, 290–4, 311–2, 318, 324
hypernonmonotonicity 85, 90, 93–6, 148, 200, 305
hysteresis 52, 91–3, 296; *see also* closure

idempotence 72, 208n20, 278n12; *see also* closure; Contraction
identity 27, 228–9, 318; generalized 173, 246–9, 275
implication xi, 4–6, 8, 10, 12–28, 36–43, 46–50, 52–3, 55–6, 58–9, 61–7, 70–2, 75–102, 103–7, 113–49, 157, 172, 179–80, 192–3, 196–200, 210–56, 281–3, 285–91, 293–307, 308, 310–14, 323–4; defeasible 21, 72–86, 93–5, 104, 114, 116–20, 132–3, 136, 193–4, 213, 239; definitions of 46–9, 52–3, 56, 106–7, 110, 179–80, 220; explicitation by 124–6; logical 58–9, 70–2, 103–31, 180, 224; material 12–28, 36–9, 46–50, 75–102, 104, 121, 143, 192, 243, 253, 300, 303, 305; *see also* consequence; implication-space semantics; incompatibility; reason relations
implication frame 212, 217, 220–1, 224–7, 229, 234, 239, 253–4, 259, 264, 266, 269, 276n7, 277n8; *see also* implicational roles; implication-space semantics
implicational role inclusion *see* role inclusion
implicational role 19, 56, 99, 210–76, 214, 221, 295–300, 310, 314; as abstractions 211–215; adjunction of 216–7, 223; bearers of 212, 230, 239; and conceptual content 215; entailment 216; implication equivalence 214; inclusion 237–53; of logically complex sentences 218, 231–2; premisory/conclusory role 23–4, 213–5; symjunction of 217–9, 223; *see also* implication-space semantics; rational forms; role inclusion
implication space xi, 18–19, 20, 22–5, 27–8, 86, 97, 199, 200, 211–12, 218, 225–9, 230–7, 290, 302, 304–6; candidate implications 19, 20, 212–4, 216–7, 228–9, 234, 242, 254, 295, 310; conclusion set 239, 242, 306; implication 53, 56, 216, 227, 253–4; implication frame of 224–6, 254, 299; premise set 239, 242, 306; *see also* implication-space semantics
implication-space semantics 19–20, 22, 24–5, 53, 56, 199–200, 210–76, 283, 286–7, 289, 293–8, 300–1, 309, 313; conceptual content in 215, 241; conic models of 242–4, 247, 299; consequence in 212, 216, 220; implicational roles 214, 222, 227; and truth-maker semantics 224–8;

see also conceptual content; implicational roles; rational forms
implicit 3, 47–9, 56, 65, 76, 83–9, 91–6, 98, 102, 147, 155, 197, 314, 323–4; *see also* explicitation/explication
impossibility 6–8, 29–31, 40, 285, 288, 291, 310, 320; as noncompossibility 6, 321; normative 164–6, 172; of states 10–12, 19, 24, 160–1, 164–6, 172, 174, 177–88, 191–4, 196–8, 224–6, 306; *see also* alethic; incompatibility, modality; truth-maker semantics
inclusion 4; content 237, 240–1; for implicational roles 238–9, 240, 243, 246, 249, 250–1, 254, 289, 299; logical 246–7, 299, 300, 303; modal 289; normative 249; rational 25, 56; role 237–46, 252; *see also* implication
incoherence 41, 94, 146, 193; defeasible 77, 155n2; persistent 93; of premise sets 77–8, 93, 115–6; *see also* incompatibility
Incoherence-Incompatibility (II) 67–68, 146, 223, 301; *see also* negation
incompatibility 3–9, 11–18, 20, 23–7, 31, 36–43, 45–53, 55–8, 60–4, 66–81, 84–7, 92–102, 103–5, 107, 115–6, 118, 122, 142–9, 155, 163, 165–72, 184, 193–4, 196–9, 206, 219, 244, 249, 276, 281–3, 285–91, 300–7, 309, 312–5, 318–26; logical 16, 38, 58, 63–4, 67–9, 70, 86, 92, 103–5, 115, 144, 146–9; material 3–9, 11–18, 20, 23–7, 36–43, 60–4, 66–81, 249, 285–91, 300–7, 312–15, 318–26; monotonicity of 78, 84–6, 104, 115–18, 146–9, 155, 193–4; symmetry of 39–42, 45, 50–53; *see also* inconsistency; negation; reason relations
inconsistency 16, 25, 36, 41, 52–3, 58–9, 62–3, 68, 75, 94–5, 97, 101n8, 149; *see also* incompatibility
indefeasible *see* persistence
inference 31, 46, 61, 65, 85–88, 106, 119, 140, 268, 288, 296; as acknowledgment of implicit commitments 89, 92, 119, 296; defeasible 86, 236; *see also* commitment; inferentialism
inference rules *see* operational rules; structural rules/principles
inferentialism 18–19, 22–4, 56–7, 61–2, 105, 123, 145, 255, 281, 294–5, 304–5; compositionaltiy of 255, 291; holism of 255, 291–2; and implicational roles 22, 56–7, 105, 145, 295; and implication-space semantics 22–4, 199, 255, 291; as theory of meaning 53, 281; *see also* implication-space semantics; inference; meaning; reason relations
interpretation 12, 19, 22, 47, 107, 295, 305, 322; for conditionals 219; in implication-space semantics 218–20, 225–6, 230, 232, 235, 241; for linear logic 232; noncontractive 230; in truth-maker semantics 175, 179, 181, 191, 192
isomorphism 180–9; between consequence relations 17, 158, 166, 170, 172–173, 178, 299; as defining higher-order representational norms 198; and metavocabularies 10, 11, 24, 27, 32, 99, 181,197, 199, 293, 305–306; between modal relations 162, 181, 184; between theories 181, 191, 195, 322; *see also* rational forms

James, William 2
justification 3–4, 35–7, 44–5, 72–3, 325n5; *see also* challenge; defense; entitlement; pragmatics

K3 *see* Strong Kleene Logic (K3)
Kant, Immanuel 83, 306–9
Kaplan, Daniel S. xi, 17–19, 23, 156n9, 276n1, 278n14, 278n17, 307n3
Kripke, Saul 1, 10, 54n2, 208, 241, 279n21

Lewis, David 1–2, 10
language 1–8, 11–12, 16, 26–7, 60–3, 65–6, 97–8, 103–4, 108, 110–12, 121–2, 126–7, 134–7, 142, 146, 154, 156–7, 169–71, 173, 175, 177, 191–3, 195–7, 217–21, 225–7, 230–2, 241–2, 250, 253–4, 266–8,

277–8, 282–9, 291–2, 302–3, 306; *see also* base; discursive practices; sentences; vocabulary
lexicon 5–7, 13–15, 19, 66, 225; extended 14–15, 27, 108, 126; *see also* vocabulary
liar sentence 190–2
Linear Logic 96, 113, 232–7, 263–6, 289, 297; connectives 236; interpretation function 232; multiplicative additive (MALL) 232–237, 263–264, 289; and phase-space semantics 234–235, 237; *see also* implication-space semantics
linguistic rationalism 3–4, 12; pragmatic 55; *see also* metalinguistic rationalism
logic *see* strict-tolerant logic (ST)
Logic of Paradox (LP) 23, 72, 209n40, 241–7, 270–2, 299; *see also* role inclusion
logical expressivism 17, 60, 60–71, 104–5, 107, 121–4, 126, 128, 139–41, 144, 147, 235, 282, 301; role of logical vocabulary 104, 122, 131, 145, 282, 301–2; *see also* conditional; expressivism; logical vocabulary; LX; negation; sentential connectives
logical vocabulary 2, 12, 18, 41, 52–4, 60–71, 76, 84–6, 91, 107–12, 120–24, 142, 147–9, 175–6, 184, 193–5, 199–200, 217–9, 223, 230, 236, 249, 253, 300, 304, 306; expressive role of 13–15, 60–71, 75, 80, 88, 94–5, 121–3, 144, 281–9, 301–2; introduction of 13–15, 97, 103–5, 107–115; as universally LX 16–17, 23, 67, 70–1, 94, 105, 107, 111–12, 121–8, 130–1, 144–7, 156n12, 199, 282, 288, 302; *see also* metavocabulary; vocabulary
logicism 57–59, 69, 86; and closure 60, 80, 85–91; and pure logic 63–64, 143; about reasons 59–60, 66, 69, 80, 82, 94, 104, 116; structural 78, 80, 82, 84–90, 92, 94, 97; *see also* logical expressivism
LX (elaborated from/explicitating of) 16–17, 64, 66, 69, 82, 96, 197; universality of 16–17, 23, 67, 70–1, 94, 105, 107, 111–2, 121, 126, 128, 130–1, 144–7, 199, 288, 302; *see also* explicitation; logical expressivism

MALL *see* Multiplicative-Additive Linear Logic
matter *see* hylomophism; rational form
McDowell, John H. 9
meaning 2–4, 6–7, 9, 11, 19–20, 22, 25–6, 29–30, 53–4, 195–7, 249–50, 283–4, 292, 295, 303–4; inferentialist 53, 56, 250; *see also* conceptual content; inferentialism; propositions; rational forms; representationalism; semantics
mereology 10, 12, 18, 196
metainference 13–14, 17, 23, 242; and conceptual roles 23, 299; validity of 242, 270–2, 299–300; *see also* validity
metalanguage 3, 98, 122, 134–5, 253; logical 65, 83, 146; mathematical 1; pragmatic 3, 26, 283; semantic 1, 3, 283
metalinguistic rationalism 4, 8, 54, 289–90, 307; *see also* linguistic rationalism
metaphysics 173, 197, 246–9, 255, 272, 291; *see also* grounding; identity
metavocabulary 3–28, 158, 195, 198, 200, 210–11, 281–94, 305–6, 319; extrinsic-explanatory 13, 199–200, 211, 281, 286–8; implication-space 18, 23–4, 27, 294, 304–5; intrinsic-explicative 13, 15, 23, 56, 199–200, 210, 228, 281, 286–8, 294, 299; logical 2, 18, 63, 99–100, 145, 288; pragmatic 2, 6–11, 34, 40, 49, 53, 55–6, 98, 158, 172, 196, 198, 281, 283–90, 302–6, 308, 319, 323; rational 6–11, 28, 53, 196, 199–200, 288, 303–6; representational 10, 210, 253, 283; semantic 2, 6–11, 34, 53, 86, 99–100, 158, 196–200, 281, 283–90, 294, 297, 304–6, 308, 323; *see also* metalanguage; vocabulary

modality 9, 30, 167, 287, 290–1, 310, 312; alethic 1, 6–7, 9, 11, 30, 167–9, 172, 181, 184, 197, 199, 228, 285, 290–4, 309, 319–4; deontic 1, 9, 11–12, 30–2, 168–9, 173, 181, 291, 309, 319, 323–4; necessity 8, 135, 148; *see also* alethic; deontic; necessity; possibility
model theory 227; Tarskian 1–2, 10, 13; *see also* implication-space semantics; model; truth-maker semantics
monoid 24, 211, 228–230, 234, 263, 294–5; elements of 228, 254, 295, 297, 309–10; idempotence 278n12; and implication frames 212, 230, 234, 254; and reason relations 228, 310
Monotonicity (MO) xii, 78–9, 82, 87–8, 101n13, 183; Cautious (CM) 82–91, 139–41; failure of xii, 72, 74, 78, 81, 85, 89–90, 93, 104, 109–11, 116–20, 131–2, 147–8, 188, 192–4; global 16, 25, 69, 76–7, 86–91, 96–9, 139–42, 147–9, 160, 182; local 76–7, 132–6; Monotonicity Box 135–136; *see also* classical logic; nonmonotonicity; weakening
Multiplicative-Additive Linear Logic 232–4, 263–4, 298
multisets 106–7, 129, 150, 155n3, 229–30, 279n19
multisuccedent 10, 22, 28, 46–7, 52, 69, 99, 145, 306; *see also* bilateralism; NonMonotonic MultiSuccedent logic (NMMS)

necessity 8–9, 135–6, 148; Persistence-Codifying Modal Operator 76, 135; *see also* implication; modality
negation 32, 67–8, 74, 116,146, 179, 297; expressive role of 68, 76, 124–5, 233, 297; and incompatibility 64, 67–8, 74, 100, 219; *see also* Incoherence-Incompatibility; incompatibility; logical vocabulary; sentential connectives
NMMS *see* NonMonotonic MultiSuccedent logic
nonmonotonicity 17, 28, 42, 73–85, 93, 95–6, 101n8, 104, 113, 116, 118, 139, 148–9, 192–4, 213, 223, 236, 289, 293, 305; hypernonmonotonicity 85, 90, 93–6, 148, 200, 305; of incompatibilities 104, 116, 193; *see also* monotonicity
NonMonotonic MultiSuccedent logic (NMMS) 17, 81–2, 92, 105, 108–115, 119–31, 134–8, 146–9, 173, 185–7, 190, 192–5, 221–4, 227–8, 236, 253–4, 292, 300–6; axioms of 108–9, 118; and implication-space semantics 221–4; rules of 109; *see also* sequent calculus
Normative Governance 159–161; *see also* representation
normative metavocabulary *see* metavocabulary
norms 38–9, 121–2, 320; for reasoning 3, 38–9, 122, 142–3, 158–9, 281–5, 288, 325n5; and representation 162–3, 167–8, 173, 184, 195–9, 209n32, 306; *see also* bilateralism; commitment; deontic; discursive practices; entitlement; pragmatics; social practices

open 90–1, 103–5, 116, 126, 147, 173, 210, 237, 255, 281, 283, 293–4, 296–7, 300–1, 303; and structure of reason relations 72, 81–2, 85, 89, 91, 96–9, 100, 148–9, 187, 189, 195, 210, 212, 305; substructurality 85, 89, 113, 120, 195; *see also* closure; reason relations
operational rules 109, 129, 183–7, 194, 201, 228; deeming impossible 184; left-rules 135, 180, 219, 222–3, 298; PN 181–3; right-rules 84, 135, 180, 219, 222, 236, 298; soundness of 136, 185, 187–9; and truth-maker consequence 178, 180–9, 191; validity of 112–3, 120, 185–7; *see also* sequent calculus
order of explanation *see* pragmatics-first; semantics-first
out of bounds 7, 10–12, 47–52, 106–7, 124–5, 156n10, 172, 181, 310, 325n7; *see also* bilateralism; norms

partition 19–20, 228–30, 254, 267–8, 270, 295, 297, 310, 312, 315; *see also* monoid
Peirce, Charles Sanders 2
Permutation 42, 110, 157n25; *see also* sequent calculus
persistence 75–8, 93–4, 97–8, 102n15, 119, 133–6, 148, 155n2, 240, 243; *see also* Monotonicity
practical attitudes 7–9, 30, 33, 53, 55; acceptance 30–38, 42–45, 55; doxastic 56, 196; expression of 8, 9, 30; rejection 30–8, 42–45; *see also* assertion; claiming; commitment; denial; speech acts
pragmatic metavocabulary *see* metavocabulary
pragmatics 2, 6–9, 30–5, 40, 99, 103, 181, 195–99, 210–11, 253, 281, 304, 307, 319; affirmation/denial and commitment/entitlement 10–11, 24, 33–5, 40–3, 46, 55–6, 99, 103, 181–4, 290, 294; discursive 2–3, 7–10,26- 40, 52, 55–57; pragmatic/semantic distinction 2–6, 24, 30, 283–89, 319; understanding 6, 10, 47–8; *see also* language; use
pragmatics-first 3, 31–2, 40–41, 55, 100, 103–4 ,142–3, 158, 173, 180–1, 207n19, 304, 308–9, 319; *see also* order of explanation
Price, Huw 2, 38, 325n5
proof-search 150–1, 202, 204, 259, 263; *see also* sequent calculus
proof theory 4, 18, 24–5, 65, 103, 278n13, 293, 298, 305–6; *see also* sequent calculus
propositions 5, 18, 42–3, 69, 100, 199, 281–3, 290–3, 309, 317; worldly 161–72, 175, 177, 179, 181–4, 195, 197, 205n3, 205n8, 210–12, 224–7, 252–3, 287, 309, 311; *see also* conceptual content; conceptual role; rational forms; representation
Putnam, Hilary 2, 207n16

Quine, Willard Van Orman 1–2, 33, 70–1, 144, 149, 207n16, 277n8

Ramsey Condition 67, 77, 79; *see also* conditional; Deduction-Detachment; explicitation; expressivism
range of subjunctive robustness 20–2, 76, 132–3, 135, 213–5, 217–9, 222–3, 229–30, 233, 235–9, 254, 257, 264, 286, 295–9; *see also* implicational roles; subjunctive robustness
rational forms 9, 12, 17, 24, 28, 159, 162–7, 181, 211, 221, 228, 253–5, 290–294, 308–324; bearers of 165, 291–2, 309–10, 314; sharing of 166, 168, 181; *see also* conceptual content; conceptual realism
rational metavocabulary 6–7, 12, 25–8, 53, 196, 284–5, 288, 290, 303–4, 306–7; extrinsic-explanatory 13, 24, 27, 199–200, 286; intrinsic-explicative 13, 15, 17–18, 23, 27, 199–200, 286, 299; *see also* metavocabulary
reasoning 2–3, 9, 26, 33, 36, 59–60, 70, 72–3, 78, 80, 82–4, 86, 90, 92, 99, 197, 303; versus reason relations 35–7, 43, 48–9; *see also* discursive practices; norms; reason relations
reason relations 3, 4–11, 33–52, 121–131, 189–195, 228–230, 283–290, 293–303; *see also* consequence; implication; incompatibility
reasons 1–6, 26, 33–46, 57–64, 69, 98–9, 101n1, 108, 125, 142–5, 172, 196, 199, 281–2, 295–6, 303–4, 322; defeasibility of 21, 72–4; for/against 21–2, 55–6, 69, 75, 100, 103, 106–7, 116, 122, 193, 289, 293; logicism 59–60, 66, 69, 78, 80–2, 87, 94–5, 104, 116, 143; material 36, 39, 52, 62–3, 66, 69, 75, 78, 80–2, 84–90, 92, 94–100, 119, 121, 180, 192–4, 291, 299, 303; question 57–9, 69; *see also* challenge; defense; inferentialism; logicism; reason relations
reasons question about logic 57–59, 69, 90; *see also* logicism
Reflexivity (RE) 82, 85, 96, 110–1, 242–3, 267, 278n18, 296, 303; *see also* Containment
rejection 30–9, 42–6, 55–6, 104, 124–5, 206n13, 283–7, 292–4, 321–2; Pragmatically Implicit (PIR)

47–8; *see also* denial; practical attitudes; speech acts
representation 7–8, 10–12, 18–20, 24–8, 31–2, 143–4, 158–71, 175, 183–4, 195–8, 210, 252–3, 283–7, 291–4, 302–7, 318, 322; and explicitation 129–31, 147, 231; factive/non-factive 161, 163, 171; and isomorphism 158–9, 168–9, 172; misrepresentation 168, 171, 209n32; and rational forms 159, 162–168, 181, 184, 195; *see also* Covariant Tracking; isomorphism; Normative Governance; rational forms
representationalism 18–19, 172, 178, 195, 210–1, 253, 281, 284–7, 304–5; *see also* inferentialism; meaning; representation
Restall, Greg 10–11, 32, 46–8, 50, 68, 106–7, 283, 293, 325n7
Ripley, David 10–11, 32, 46–8, 50, 101n3, 190–1, 202, 209n38, 241
role inclusion 237–55, 299–300, 303; definition of 238–9; and strong Kleene scheme 214–46; and metaphysics 246–9; nonlogical 249–53; *see also* implicational role; Logic of Paradox (LP); strict-tolerant logic (ST); Strong Kleene Logic (K3); tolerant-strict logic (TS)
Rorty, Richard 2
Russell, Bertrand 1, 58–60

self-consciousness, rational 25, 28, 69, 122–123, 126, 294, 306–307; conceptual 28, 69, 122, 126, 306–7; and logical vocabulary 69, 122–3; *see also* consciousness; metavocabulary
Sellars, Wilfrid 2, 36, 62, 100, 164, 287; *see also* inferentialism
semantic metavocabulary *see* metavocabulary
semantics 1–4, 30–2, 39, 43, 83, 99, 143, 190, 195, 206n10, 233, 255, 305; hyperintensional 2, 173; inferentialist 18, 22–4, 53, 255, 294–5; intensional 2, 10, 35, 76–7; intrinsic 19, 22, 24; possible worlds 2, 174, 216, 326; representionalist 24, 294; *see also* implication-space semantics; truth-maker semantics
semantics-first 3, 158–62, 172–3, 308–9, 319; *see also* order of explanation
sentential connectives 14–17, 63–70, 76, 81, 103, 129, 145–9, 218, 228, 231–3, 236–7, 255, 298; *see also* conditional; conjunction; disjunction; negation
sequent 105, 107–12; *see also* sequent calculus
sequent calculus 13–18, 22–4, 41, 68–9, 103–12, 134–5, 141–9, 173, 178, 180–1, 185, 189–90, 194, 211, 221–3, 227, 230–1, 236, 249, 288–9, 298, 300–1, 304–5; axioms of 104–5, 108–9, 118, 300; reduction trees 202; rules of 106–12, 181–9, 222–3, 231, 250–2, 298; *see also* inferentialism; logical expressivism; operational rules
Sharing Rational Forms 166, 181
sign design 5, 100, 164, 167
Simonelli, Ryan xii, 46–7, 50, 53, 79, 156n17
social practices 1–2, 197; *see also* discursive practices; language
speech acts 3, 9, 30, 52, 55–6, 196; *see also* assertion; denial; metavocabulary; practical attitudes
Spinoza, Baruch 11, 59, 291
ST *see* strict-tolerant logic (ST)
Stalnaker, Robert 2
state of affairs 6, 160; *see also* fact; worldly state
Strong Kleene Logic (K3) 23, 72, 209n40, 241–4, 246–7, 270–2, 299; *see also* Logic of Paradox (LP); role inclusion; strict-tolerant logic (ST); tolerant-strict logic (TS)
strict-tolerant logic (ST) 23, 86, 190–2, 241–3, 255, 279n22, 289, 299
structural rules/principles 15, 41, 78, 83, 87, 89, 110–3, 157n25, 298, 326n7; and NMMS 110–3, 116, 131–2, 139–40, 142–3, 147–8, 149; and truth-maker consequence 178, 180, 182, 187–9; and implication-space semantics 239, 278–9; *see also* Containment; Cut;

NonMonotonic MultiSuccedent logic (NMMS); sequent calculus; Weakening
structure *see* structural rules
subjunctive conditionals 34–5, 73–4, 291; *see also* subjunctive robustness
subjunctive robustness 9, 98, 292; *see also* range of subjunctive robustness
succedent 10, 17, 20, 39, 46–7, 52, 96, 108, 115; *see also* sequent calculus
Succedent-Summation (SS) 115, 146, 223, 301
symmetry 39–41, 45, 49–50, 168, 171, 225–6; *see also* incompatibility

Tarski, Alfred 1–2, 10, 25, 41, 47, 72, 81, 83, 87–8, 97, 101n6, 102n14, 149, 305
theorems 81, 86, 118–20, 136, 247
tolerant-strict logic (TS) 23, 209n40, 241–3, 246, 255, 270–2, 299
tonk 65
Tractatus Logico-Philosophicus 1, 59
transitivity 16–17, 23, 70, 78–85, 95–8, 113, 116, 147, 187, 190–2, 293, 296, 305; local 139–42; *see also* closure; Cumulative Transitivity (CT); Cut
Travis, Charles 2
truth 19, 29–35, 53, 55, 57–8, 77, 116, 159–64, 173–4, 190–2, 220, 292–5, 317–19, 325n5–7; *see also* truth-makers; truth-values
truth conditions 20–1, 34–5, 66, 173, 207n16, 295; *see also* truth-maker; worldly states
truth-maker 159–62, 173–176, 225–6; of sets 176
truth-maker semantics 24, 53, 187, 192, 194, 197, 293–4, 306; Consequence 179–80, 184–9; Containment 178, 187–9; Downward-Closure 177, 187–9; Entailment 178; Exclusivity 177, 187–9; Exhaustivity 177, 187–9; fusion 174; and implication-space semantics 224–8; and logical consequence 180; logical vocabulary 175–6; modalized state space 174; models of 175; null state 175; parallel models 226, 259; Propositional Implication 179
truth-predicate 190–1, 195, 209n36, 242
truth-values 19–20, 29–35, 42, 173, 241–2, 244–5, 325n7; *see also* bivalence; falsity; truth; truth-maker semantics
turnstile 10, 39, 43, 46–7, 50, 76–7, 84, 87–9, 139, 146, 156n16; *see also* implication

use 2–4, 7–8, 26–27, 30, 52–54, 63–64, 94, 111–12, 123, 195–8, 283–4, 292; and representation 8, 11–12, 32, 143, 158–9, 162–4, 168–70, 173, 197–198, 209n32, 287, 306; *see also* assertion; denial; norms; pragmatics; speech acts

validity 58, 69, 84, 94, 131, 157n21, 178, 180, 184–9, 201–4, 208n26, 242; classical 112–5, 119–20, 194; metainferential 270–2, 299–300; *see also* consequence
vocabulary 5–28, 103–5; autonomous 26, 105–112, 126, 131, 196; base 5–28, 61–71, 75–6, 81–2, 92–7, 107–14, 126–31, 135, 144–9, 193–6, 199–200, 221–2, 231, 286, 300–5; logical 12, 41, 60–71, 107–12, 121–24, 175–6, 218–9; logically extended 14–17, 27, 63–6, 75–6, 81, 95–8, 107–13, 118–21, 126–31, 137, 144–9, 222, 231, 300–2; nonlogical 13–16, 39, 58–9, 62–8, 70–1, 92–5, 105, 108, 128, 144, 149, 300; *see also* lexicon; metavocabulary; reason relations; sentences

Weakening 79, 110–1, 119, 133–4, 151, 183, 187–9, 201, 298; *see also* closure; Monotonicity; sequent calculus; structural rules
Whitehead, Alfred N. 60
Wittgenstein, Ludwig 1–2, 9, 59, 100, 160–1
worldly states 9, 24–5, 32–4, 164–7, 196–9, 228, 253, 283–8, 290–5, 308–10, 322; *see also* fact; fusion